


NCSC-TG-028
VERSION-1



NATIONAL COMPUTER SECURITY CENTER

ASSESSING CONTROLLED ACCESS PROTECTION

19980309 294

PLEASE RETURN TO:

25 May 1992

BMD TECHNICAL INFORMATION CENTER
BALLISTIC MISSILE DEFENSE ORGANIZATION
7100 DEFENSE PENTAGON
WASHINGTON D.C. 20301-7100

Approved for Public Release:
Distribution Unlimited

EXCLUDED FROM AUTOMATIC DOWNGRADING AND DECLASSIFICATION

U3738

Accession Number: 3738

Publication Date: May 25, 1992

Title: Assessing Controlled Access Protection

Corporate Author Or Publisher: National Computer Security Center, 9000 Savage Rd., Ft. GG Meade, MD 2 Report Number: NCSC-TG-028

Descriptors, Keywords: Rainbow Technical Guideline Controlled Access Computer Evaluation Criteria Protection AIS

Pages: 00069

Cataloged Date: Sep 15, 1992

Document Type: HC

Number of Copies In Library: 000001

Record ID: 24734

FOREWORD

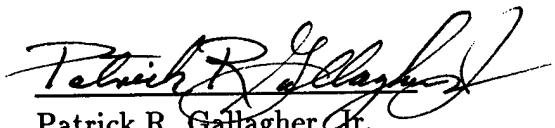
The National Computer Security Center is publishing *Assessing Controlled Access Protection* as part of the "Rainbow Series" of documents our Technical Guidelines Program produces. In the Rainbow Series, we discuss in detail the features of the *Department of Defense Trusted Computer System Evaluation Criteria* (DoD 5200.28-STD) and provide guidance for meeting each requirement. The National Computer Security Center, through its Trusted Product Evaluation Program, evaluates the security features of commercially-produced computer systems. Together, these programs ensure that organizations are capable of protecting their important data with trusted computer systems.

Assessing Controlled Access Protection explains the controlled access protection requirements of the *Trusted Computer System Evaluation Criteria*. The guide's target audience is the technical analysts tasked by the Department of Defense components to determine whether a system meets these requirements.

As the Director, National Computer Security Center, I invite your recommendations for revision to this technical guideline. We plan to review and update this document periodically in response to the needs of the community. Please address any proposals for revision through appropriate channels to:

National Computer Security Center
9800 Savage Road
Ft. George G. Meade, MD 20755-6000

Attention: Chief, Standards, Criteria, and Guidelines Division


Patrick R. Gallagher, Jr.
Director
National Computer Security Center

May 1992

ACKNOWLEDGMENTS

The National Computer Security Center expresses appreciation to Dr. Dixie B. Baker, of The Aerospace Corporation, as the principal author of this document, and Ms. Caralyn Crescenzi as project manager.

We also thank the evaluators, vendors, and users in the United States computer security community who contributed their time and expertise to the review of this document.

Executive Summary

Assessing Controlled Access Protection provides guidance to the Department of Defense components charged with ensuring that the automated information systems (AISs) used for processing sensitive or classified information provide at least controlled access protection.

The objectives of this guideline and its supporting documentation set are:

1. To provide a methodology for performing a technical analysis to support the certification of controlled access protection in AISs submitted for accreditation;
2. To provide an interim approach for achieving controlled access protection until a suitable NSA-evaluated product is available; and
3. To clarify the intent, security functionality, and level of assured protection that controlled access protection provides.

The guidance provided in this document is targeted toward multi-user AISs designed for DoD operations in system-high security mode and in dedicated mode, where directed by the DAA. This guidance does not specifically address connectivity with a local-area or wide-area network. Nor does it address related areas such as physical security, TEMPEST, communications security, or administrative security (e.g., trusted distribution).

This guideline is written to serve as the synergist that integrates and consolidates information contained in the following documents into a unified explanation of the requirements for and intent of controlled access protection.

- *A Guide to Understanding Audit in Trusted Systems*
- *A Guide to Understanding Configuration Management in Trusted Systems*
- *A Guide to Understanding Design Documentation in Trusted Systems*
- *A Guide to Understanding Discretionary Access Control in Trusted Systems*
- *A Guide to Understanding Identification and Authentication in Trusted Systems*
- *A Guide to Understanding Object Reuse in Trusted Systems*
- *A Guide to Writing the Security Features User's Guide for Trusted Systems*
- *Guidelines for Writing Trusted Facility Manuals*

- *Trusted Product Evaluation Questionnaire*

The National Computer Security Center (NCSC) publishes and distributes these documents to support the certification and accreditation of AISs required to provide controlled access protection. To request copies of these documents, contact the National Technical Information Service (NTIS).

Contents

1	BACKGROUND	1
1.1	NATIONAL POLICY	1
1.2	SECURITY ACCREDITATION	2
1.3	TRUSTED PRODUCT EVALUATION	3
1.4	SCOPE AND PURPOSE	5
2	CONTROLLED ACCESS PROTECTION	9
3	ARCHITECTURAL FOUNDATION	13
3.1	TRUSTED COMPUTING BASE	13
3.2	ENFORCEMENT	17
3.3	DOMAIN SEPARATION	18
3.4	DEFINED SUBSET	20
3.5	RESOURCE ISOLATION	20
4	PROTECTION MECHANISMS	22
4.1	IDENTIFICATION & AUTHENTICATION	22
4.2	DISCRETIONARY ACCESS CONTROL	24
4.3	OBJECT REUSE	28
4.4	AUDIT	29
5	DOCUMENTATION AND LIFE-CYCLE ASSURANCE	33

ASSESSING CONTROLLED ACCESS PROTECTION	viii
5.1 DESIGN DOCUMENTATION	33
5.2 SYSTEM INTEGRITY	34
5.3 CONFIGURATION MANAGEMENT	35
5.4 TRUSTED FACILITY MANUAL	37
5.5 SECURITY FEATURES USER'S GUIDE	38
5.6 TESTING	39
6 TECHNICAL ANALYSIS	41
6.1 SELECTION OF ANALYSTS	41
6.2 TECHNICAL ANALYSIS PROCESS	42
7 RISK MANAGEMENT	53
7.1 PROTECTION LIMITATIONS	54
7.2 IDENTIFIED DEFICIENCIES	55
7.2.1 SYSTEM ARCHITECTURE	55
7.2.2 IDENTIFICATION AND AUTHENTICATION	56
7.2.3 DISCRETIONARY ACCESS CONTROL	56
7.2.4 OBJECT REUSE	56
7.2.5 AUDIT	56
7.2.6 SYSTEM INTEGRITY	57
8 ACRONYMS	63
9 GLOSSARY	65

List of Figures

1.1	National Policy on Controlled Access Protection	1
1.2	DoDD 5200.28 Timetable for C2	2
3.1	Trust Hierarchy in an AIS	13
3.2	Relationship between System Engineering and Assurance	16
3.3	TCSEC C2 System Architecture Criterion	17
4.1	TCSEC C2 Identification and Authentication Criterion	23
4.2	TCSEC C2 Discretionary Access Control Criterion	24
4.3	ACL for File <i>georges_data</i>	26
4.4	Output from Directory Study	27
4.5	Unix Command Sequence	27
4.6	TCSEC C2 Object Reuse Criterion	28
4.7	TCSEC C2 Audit Criterion	30
5.1	TCSEC C2 Design Documentation Criterion	33
5.2	TCSEC C2 System Integrity Criterion	35
5.3	TCSEC C2 Trusted Facility Manual Criterion	37
5.4	TCSEC C2 Security Features User's Guide Criterion	38
5.5	TCSEC C2 System Testing Criterion	39
6.1	Controlled Access Protection Technical Analysis Process	43

List of Tables

2.1	Security Policy Control Objectives and Implementation Requirements	11
4.1	Object Reuse Mechanisms	29

Chapter 1

BACKGROUND

1.1 NATIONAL POLICY

In July of 1987, the Federal government issued the National Policy on Controlled Access Protection [36], establishing the policy for automated information systems (AISs) that are accessed by multiple users with different authorizations to the information contained in the system. The Policy, shown in Figure 1.1, mandates that these systems provide automated controlled access protection and that this minimal level of protection be provided within five years of the Policy's issuance. The Policy gives the Federal agencies responsibility for ensuring that its provisions are carried out.

All automated information systems that are accessed by more than one user, when those users do not have the same authorization to use all of the classified or sensitive unclassified information processed or maintained by the automated information system, shall provide automated Controlled Access Protection for all classified and sensitive unclassified information. This minimum level of protection shall be provided within five years of the promulgation of this policy.

Figure 1.1: National Policy on Controlled Access Protection

The Department of Defense (DoD) carries the Policy forward in Directive 5200.28, *Security Requirements for Automated Information Systems (AISs)* [38], which specifies requirements for AISs that handle classified, sensitive unclassified, or unclassified information. The Directive provides a risk-assessment procedure, extracted from CSC-STD-003-85 [11], which is used to determine the minimum *Trusted Computer System Evaluation Criteria (TCSEC)* [14] evaluation class required for an AIS, based on the sensitivity of the information stored in or processed by the AIS and on the clearances of its users. For AISs that process or handle classified and/or sensitive un-

classified information, and that, based upon the prescribed risk-assessment procedure, require at least controlled access protection, the Directive mandates an implementation timetable of 1992, as shown in Figure 1.2.

All AISs that process or handle classified and/or sensitive unclassified information and that require at least controlled access protection (i.e., class C2 security), based on the risk assessment procedure described in enclosure 4, shall implement required security features by 1992.

Figure 1.2: DoDD 5200.28 Timetable for C2

The National Security Agency (NSA) evaluates commercial products designed to meet the TCSEC requirements and lists them in its Evaluated Products List (EPL) [34] maintained by the National Computer Security Center (NCSC). The Directive tasks the NSA to serve as a focal point for technical matters relating to the use of trusted computer products and to provide to the Department of Defense (DoD) components, as requested, technical assistance in evaluating and certifying computer-based security features of AISs used in operational environments. This guideline is responsive to this tasking; its purpose is to provide the DoD components technical guidance to support the certification and accreditation of operational systems.

1.2 SECURITY ACCREDITATION

Prior to allowing an AIS to handle any classified or sensitive information, a Designated Approving Authority (DAA) must accredit it to operate in one of three security modes: dedicated, system high, or multilevel. In dedicated mode, all users have the clearance or authorization and a need-to-know for all data handled by the AIS. In system high mode, all users have a security clearance or authorization, but not necessarily a need-to-know, for all data handled by the AIS. Multilevel mode allows two or more classification levels to be processed simultaneously within the same AIS when not all users have a clearance or formal access approval for all data handled by the AIS.

A program for conducting periodic review of the adequacy of the safeguards for operational, accredited AISs also must be established. [38] The DAA should be involved in all phases of the system acquisition, beginning with the development of the security policy and operations concept, and including the specification of the security requirements, reviews conducted during the design and development phases, and security testing, to ensure that he or she understands the operational needs, how system components work together, how the system interfaces with other systems and organizations, and the risks associated with the system.

The technical evaluation of an AIS's security features and other safeguards, made in support of the accreditation process, is called *certification*. Certification establishes the extent to which a particular AIS's design and implementation meet a set of specified security requirements. *Accreditation* is the DAA's formal declaration that an AIS is approved to operate in a particular security mode, using a prescribed set of safeguards. Accreditation is the official management authorization for operation of an AIS and is based on the certification process as well as other management considerations. The accreditation statement affixes security responsibility with the DAA and shows that due care has been taken for security. [38] Although certification involves a great deal more than the technical analysis described in this document, the guidance contained herein can provide a technical basis for the certification portion of the accreditation process.

1.3 TRUSTED PRODUCT EVALUATION

The DoD policy specified in DoDD 5200.28 states that:

Computer security features of commercially produced products and Government-developed or -derived products shall be evaluated (as requested) for designation as trusted computer products for inclusion on the Evaluated Products List (EPL). Evaluated products shall be designated as meeting security criteria maintained by the National Computer Security Center (NCSC) at NSA defined by the security division, class, and feature (e.g., B, B1, access control) described in DoD 5200.28-STD.

The NCSC maintains the EPL and, using technical support from NSA, evaluates, assigns ratings to, and enters onto the EPL products designed and developed in accordance with the TCSEC. NSA maintains a cadre of trusted-product evaluators both from within the agency and from Federally Funded Research and Development Corporations (FFRDCs). The trusted product evaluation program (TPEP), described in detail in *Trusted Product Evaluations: A Guide for Vendors* [41], comprises the following five phases:

1. *Proposal Review*. When a vendor requests that its product be evaluated for possible inclusion on the EPL, NSA prescreens the proposed product relative to its usefulness to DoD components, its technical merit (through an intensive Preliminary Technical Review), and the vendor's commitment to the product.
2. *Vendor Assistance*. If NSA decides that the product has potential merit, it signs a Memorandum of Understanding (MOU) with the vendor. Through this MOU, the vendor agrees (among other things) to give NSA evaluators access to

the highly proprietary hardware and software design documentation needed to perform an evaluation. Once the MOU is signed, NSA assigns a small evaluation team to track the product through its development and to provide assistance in the interpretation and application of TCSEC requirements for the targeted class. This team works closely with the vendor throughout the development of the product to help determine the targeted division and class and to ensure that the design and developmental approach are compliant with the requirements of the TCSEC for that class.

3. *Design Analysis.* When development is complete, and all of the required documentation is nearing completion, the product enters Design Analysis. During this phase, an expanded evaluation team completes training (to the level of an applications programmer, for systems targeted for up to class B1, and to the level of a system programmer, for systems targeted for the higher classes). The team analyzes the product relative to the TCSEC requirements and writes a detailed Initial Product Assessment Report (IPAR). For products targeted at B2 and above, a preliminary architecture study is conducted, and at A1, the team begins examining the formal verification during this phase. Information necessary for design analysis is gained through thorough review of the hardware and software design documentation, examination of drafts of TCSEC-required documentation (e.g., Security Features Users' Guide, Trusted Facility Manual, test plans and procedures), and interactions with the vendor. Because both team members and vendor personnel are likely to be widely dispersed geographically, electronic communications are relied upon heavily for team and vendor communications. Once the analysis is completed, the team presents the IPAR to NSA's Technical Review Board (TRB), which serves as one of the TPEP's primary quality-control mechanisms. Based upon the IPAR and the team's presentation, the TRB provides to NSA management a recommendation as to whether the product is ready to begin the Evaluation Phase.
4. *Evaluation.* This phase is the actual security evaluation of the product. During this phase, the evaluation team completes the design analysis, building upon the information contained in the IPAR. Prior to beginning functional testing, the team presents its assessment to the TRB, with a request that the evaluation be allowed to proceed to testing. The team then conducts functional testing (all classes) and penetration testing (class B2 and above), examines the final versions of required documentation, and completes the Final Evaluation Report. At class B2 and above, a system architecture study and covert channel analysis are conducted, and at A1, the formal verification is validated. At the end of this phase, the evaluation team again appears before the TRB to present its findings and to recommend a final rating. Successful completion of this phase results in placement of the vendor's product on the EPL.

5. *Rating Maintenance.* NSA's RAting Maintenance Phase (RAMP) provides a mechanism for ensuring the continuing validity of a rating extended to successive versions of the rated product.

The EPL, published semi-annually as part of the *Information Systems Security Products and Services Catalogue* and updated quarterly,¹ provides system acquisition agents a good selection of C2-rated products from which to select platforms for their applications. In addition, the EPL contains a number of products that have been rated B1 and above; all of these contain acceptable controlled access protection mechanisms and, if appropriately configured, could be used in a system-high or dedicated environment. In fact, some system-high environments, particularly those with external interfaces to systems at different levels, might benefit from the additional labeling capability that Divisions B and A systems provide. Further, more and more computer vendors are bringing their products to the NSA with the request that they be considered for evaluation.² This being the case, a reasonable expectation is that the EPL will continue to expand as more vendors recognize the commercial value of NSA-rated products.

However, an assessment methodology and trained analysts are needed for those DoD programs for which a suitable NSA-rated C2 (or above) product does not exist or that do not currently have the resources necessary to rehost their software on a rated product. This guideline addresses these needs.

1.4 SCOPE AND PURPOSE

This document is intended to be used by individuals tasked to perform a technical analysis of an AIS in support of its certification and accreditation. The distinction between the terms "automated information system" and "trusted product" is important in this context. As defined in the Directive, an *automated information system* is *any assembly of computer hardware, software, and/or firmware configured to collect, create, communicate, compute, disseminate, process, store, and/or control data or information.* [38] In this guideline, the term "AIS" (or "system") refers to an AIS that is configured for a specific purpose relevant to the DoD component for which it is being accredited. The Directive defines a *trusted product* as *a product that has been evaluated and approved for inclusion on the Evaluated Products List (EPL).*[38] An AIS may be built on a trusted product (or "EPL product").

¹To obtain a copy of the current EPL, write to the National Technical Information Service (NTIS), 5285 Port Royal Road, Springfield, VA 22161.

²See Potential Products List in the *Information Systems Security Products and Services Catalogue*.

This guideline serves to unify, interpret, and apply information contained in other documents published by the NCSC. The following documents are incorporated by reference to support the technical analysis of controlled access protection.

- *A Guide to Understanding Audit in Trusted Systems* discusses issues involved in implementing and evaluating an audit mechanism. It provides guidance to vendors on how to design and incorporate effective audit mechanisms into their systems, and it contains guidance to implementors on how to make effective use of the audit capabilities that trusted systems provide. [1]
- *A Guide to Understanding Configuration Management in Trusted Systems* provides guidance to developers of trusted systems on what configuration management is and how it may be implemented in the system's development and life cycle. It stresses the importance of configuration management for all systems and suggests how it can be implemented. [2]
- *A Guide to Understanding Design Documentation in Trusted Systems* provides guidance in understanding and meeting the TCSEC's design documentation requirements. It stresses the importance of good design documentation in maintaining security throughout a system's life cycle and describes the design documentation necessary to support product review and evaluation. [4]
- *A Guide to Understanding Discretionary Access Control in Trusted Systems* discusses issues involved in designing, implementing, and evaluating discretionary access control (DAC) mechanisms. [5]
- *A Guide to Understanding Identification and Authentication in Trusted Systems* describes the identification and authentication (I&A) requirements and provides guidance to vendors on how to design and incorporate effective I&A mechanisms into their systems. [6]
- *A Guide to Understanding Object Reuse in Trusted Systems* describes the object reuse requirement and provides guidance to vendors on how to design and incorporate effective object reuse mechanisms into their systems. [7]
- *A Guide to Writing the Security Features User's Guide for Trusted Systems* explains the motivation and meaning of the TCSEC requirement for a Security Features Users' Guide (SFUG) in terms of audience, content, and organization. It is addressed to potential SFUG authors. [8]
- *Guidelines for Writing Trusted Facility Manuals* presents issues involved in writing a Trusted Facility Manual (TFM). It provides guidance to vendors on how to document functions of trusted facility management and recommends structure, format, and content to satisfy the TCSEC requirements. [32]

- *Trusted Product Evaluation Questionnaire* contains a list of questions that address the TCSEC criteria from class C1 through A1. It was developed to serve as a tool for formalizing the data-gathering process required during various phases of the TPEP. [40]

The objectives of this guideline and its supporting documentation set are:

- To provide a methodology for performing a technical analysis to support the certification of controlled access protection in AISs submitted for accreditation.
- To provide an interim approach for achieving controlled access protection until a suitable NSA-evaluated product is available.
- To clarify the intent, security functionality, and level of assured protection that controlled access protection provides.

The results of this analysis also can provide valuable information to system developers and integrators attempting to compose components into complex systems. In composed systems (e.g., networks), this assessment will provide assurance that each individual AIS provides the required level of controlled access protection. Thus this analysis will be useful in conducting an evaluation by parts [39] of the total system.

The guidance provided in this document is targeted toward multi-user AISs designed for DoD operations in system-high security mode and in dedicated mode, where directed by the DAA. This guidance does not specifically address connectivity with a local-area or wide-area network. Nor does it address related areas such as physical security, TEMPEST, communications security, or administrative security (e.g., trusted distribution).

This guide's primary audience is the analysts tasked to perform a technical assessment of an AIS's controlled access protection features and assurances. The analyst should begin by reading Chapter 2, which defines the security policies enforced by controlled access protection and explains how the requirements are derived from these policies. The analyst then should review Chapter 3, which discusses the architectural foundation necessary for controlled access protection, and Chapter 4, which describes the security mechanisms that are built upon it. A good understanding of the information contained in Chapters 3 and 4 is critical to the technical analysis process.

To gain an understanding of the documentation required as evidence that the system was built securely and that it can be operated and maintained without jeopardizing its inherent security, the analyst should next review Chapter 5, which addresses life-cycle assurances. Building upon the information contained in these chapters, Chapter 6 describes a process for performing a technical analysis to determine whether an AIS provides adequate controlled access protection. This analysis is intended to serve as

the technical basis for certification to support system accreditation. Any security analysis involves a trade-off between provided protection and assumed risk. Finally, Chapter 7 discusses risk management and identifies risks that controlled access protection is incapable of countering and risks resulting from deficiencies which may be identified during the technical analysis. Important terms are *italicized* in the text and defined in the Glossary (Appendix 9).

Chapter 2

CONTROLLED ACCESS PROTECTION

AIS security is concerned with controlling the way in which an AIS can be used; that is, controlling how users can access and manipulate the information it processes. Deriving the security requirements for a given AIS requires precise definition of the objectives of the desired control; i.e., the system's *security policy*. These control objectives will vary depending upon the perceived threats, risks, and goals of the organization for which the AIS is being accredited. Controlled access protection (as defined in the TCSEC) is founded on objectives relating to three basic types of control: security policy enforcement, accountability, and assurance. All of the requirements for AISs providing controlled access protection are derived from these objectives [14], as shown in Table 2.1 on page 11.

Controlled access protection policies are based upon a fundamental assumption that the AIS processing environment is one of mutually trusting and cooperating users. Recognition of this fact is critical to understanding the objectives of controlled access protection. The features, assurances, and most importantly the underlying system architecture of an AIS that provides controlled access protection are not intended and do not purport to prevent malicious or concerted actions aimed at circumventing the protection provided.

Controlled access protection asserts that the AIS provides:

- Protection and control over who can logon to the system.
- Mechanisms that will enable the AIS to make decisions regarding access to resources based upon the expressed wishes of its users (with no assurance that concerted, malicious actions cannot circumvent this mechanism).
- The capability to generate a reliable log of user actions and to guarantee its

correctness.

Controlled access protection is sufficient for AISs operating in system-high or dedicated security modes. However, if the AIS exports classified information that requires assured classification labeling or information that is sent to a dedicated or system-high AIS at a lower classification level, controlled access protection is not sufficient.¹ Adequate treatment of these cases is beyond the scope of this guidance.

¹Some AIS environments with integrity concerns may enforce a policy that prohibits exportation to higher levels as well.

Control Objectives	Derived Requirements
<p>Security Policy: A statement of intent with regard to control over access to and dissemination of information, to be known as the security policy, must be precisely defined and implemented for each system that is used to process sensitive information. The security policy must accurately reflect the laws, regulations, and general policies from which it is derived.</p> <p>Discretionary Security: Security policies defined for systems that are used to process classified or other sensitive information must include provisions for the enforcement of discretionary access control rules. That is, they must include a consistent set of rules for controlling and limiting access based on identified individuals who have been determined to have a need-to-know for the information.</p>	<p>System Security Policy</p> <p>Discretionary Access Control Object Reuse</p>
<p>Accountability: Systems that are used to process or handle classified or other sensitive information must assure individual accountability whenever a discretionary security policy is invoked. Furthermore, to assure accountability the capability must exist for an authorized and competent agent to access and evaluate accountability information by a secure means, within a reasonable amount of time, and without undue difficulty.</p>	<p>Identification and Authentication Audit</p>
<p>Assurance: Systems that are used to process or handle classified or other sensitive information must be designed to guarantee correct and accurate interpretation of the security policy and must not distort the intent of that policy. Assurance must be provided that correct implementation and operation of the policy exists throughout the system's life-cycle.</p>	<p>System Architecture System Integrity Security Testing Configuration Management Design Documentation Trusted Facility Manual Security Features User's Guide</p>

Table 2.1: Security Policy Control Objectives and Implementation Requirements

Chapter 3

ARCHITECTURAL FOUNDATION

Computer system architecture is the foundation upon which all AIS trustworthiness is built. This chapter discusses system architecture as it relates to trust and the concept of a Trusted Computing Base.

3.1 TRUSTED COMPUTING BASE

Inherent in the concept of trust is some assurance that the trusted person or entity possesses the required strength, capability, and integrity to merit that trust. In the case of AISs, trust is built from the bottom (i.e., hardware) up, with each layer “trusting” its underlying layer to perform the expected services in a reliable and trustworthy manner, as shown in Figure 3.1.




Figure 3.1: Trust Hierarchy in an AIS

Each layer trusts all of its underlying layers to reliably provide the expected services

and behavior. The users trust the applications they run to behave in the manner they expect; the application trusts the system calls it makes to the operating system to produce the documented results; and the operating system trusts the hardware to behave in a consistent and safe manner. Note that trust is meaningful only relative to the behaviors and strengths expected; for example, the application layer cannot expect the operating system to detect all bugs in user programs. This is particularly important relative to the trust implied for controlled access protection.

This trust hierarchy is the basis for the concept of a *Trusted Computing Base* (TCB) that cannot be compromised from above and that is always invoked to enforce a security policy with some degree of assurance. For any AIS, the TCB includes all of the software, firmware, and hardware components responsible for enforcing the security policy and all components capable of affecting the correct operation of the security mechanisms (see Chapter 4). Thus the TCB includes components whose job is to perform some function required to enforce the security policy (e.g., programs that check access-control settings on files) and components that have no direct functionality relative to the security policy, but require the capability to violate some part of the security policy of the system (i.e., *privilege*) in order to operate and therefore must be trusted (e.g., an I/O driver).

The TCSEC asserts that a trusted system architecture must exhibit protection properties that will enforce this trust hierarchy. Thus the concept of a reference monitor (or reference validation mechanism) is introduced. The term *reference monitor* represents an abstraction of the portion of the TCB that actually validates references to objects and grants (or denies) access to them. Among the properties that the reference monitor should exhibit are that it be noncircumventable (i.e., always invoked), tamperproof, and small enough to be analyzed and tested. The TCSEC imposes increasingly strict architectural and system engineering requirements on the TCB at higher and higher classes of trustworthiness. As shown in Figure 3.2, the more system engineering goes into designing the TCB, the more assured is the trust that it provides. In this figure, the increasing system engineering requirements are shown in italics beside each conceptual machine class. For classes C2 and B1, the reference monitor need not be differentiated from the rest of the TCB (which comprises the entire operating system), so that applications must trust essentially all of the operating system and hardware. Class B2 requires more system engineering to ensure that the TCB comprises largely independent modules, thus producing an additional layer of trust, as the TCB is isolated from non-security-relevant operating-system services. Classes B3 and A1 system architectures provide layered protection, with all layers ultimately reliant upon a small, conceptually simple, tamperproof, and noncompromisable reference monitor that plays a central role in enforcing the internal structuring of the TCB and the system. As the illustration shows, applications running on a class-C2 AIS (i.e., one designed to provide only controlled access protection) must trust the entire operating system and all of the hardware (i.e., all

physical resources) and firmware upon which it depends.

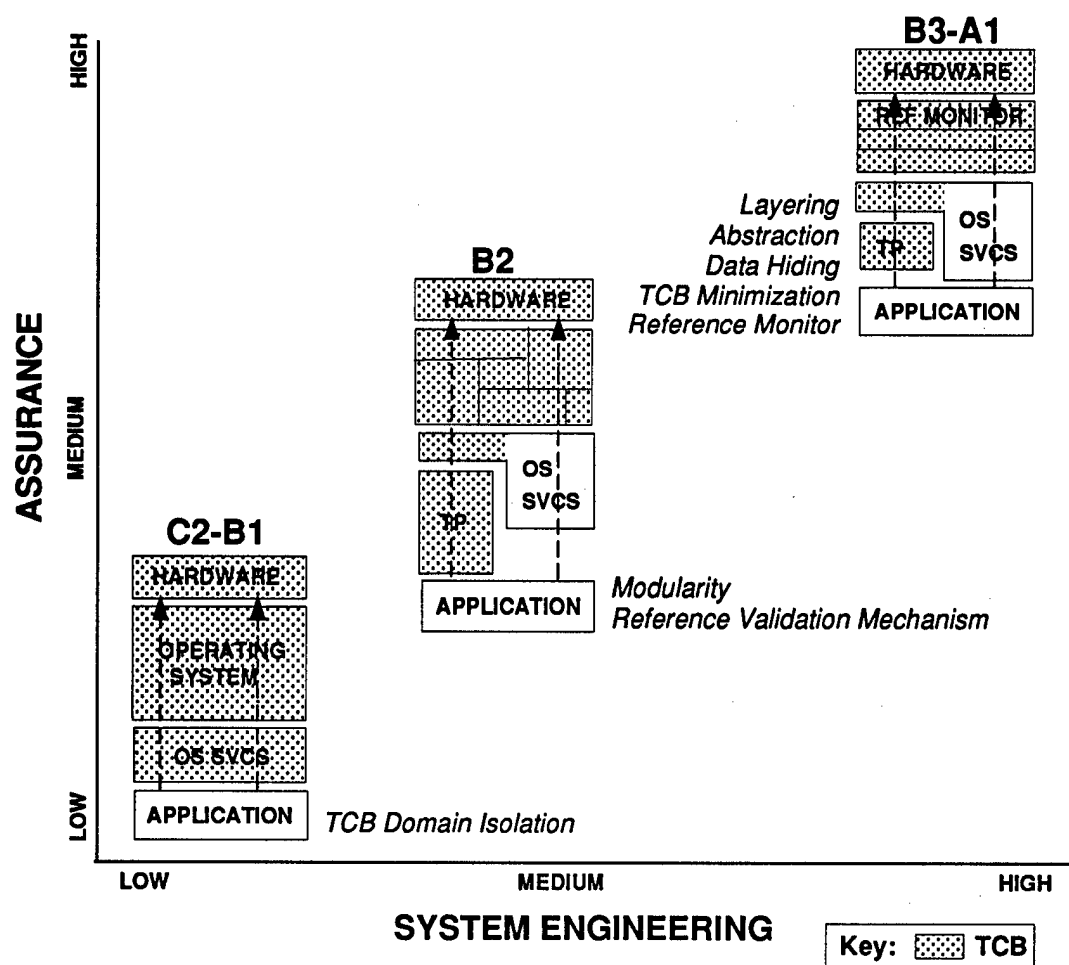


Figure 3.2: Relationship between System Engineering and Assurance

The objective and result of the TCSEC's conceptual hierarchy of trust are that demonstrating assurance in the trustworthiness of the TCB becomes increasingly tractable and assured as one progresses up the TCSEC hierarchy of trust. At class C2, the TCB may be large, dispersed, and generally unstructured; as a result, it presents a great challenge to both evaluators and persons responsible for maintaining the system's security. At class B2, the TCB still may be large, but the fact that it is modular and the result of sound software engineering practices makes it easier to understand, evaluate, and maintain than lower-rated products; thus, added assurance in its trustworthiness results. At classes B3 and A1, the TCB is small, layered, and highly structured, thus lending itself to rigorous analysis and testing, and to formal verification (A1).

3.2 ENFORCEMENT

Assurance of trust requires enforcement of the AIS's security policy. "Enforcement" implies consistency, reliability, and effectiveness. In order for a TCB to enforce the security policy, it must be both tamperproof and noncompromisable. The System Architecture criterion shown in Figure 3.3 addresses these attributes.

TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). Resources controlled by the TCB may be a defined subset of the subjects and objects in the ADP system. The TCB shall isolate the resources to be protected so they are subject to the access control and auditing requirements.

Figure 3.3: TCSEC C2 System Architecture Criterion

The term *object* refers to any passive entity that contains or receives information (e.g., files, directories, records, blocks, pages, segments, programs, video displays, printers), and access to an object implies access to the information it contains. A *subject* is any active entity in the system (e.g., person, process, device) that causes information to flow among objects or changes the system state (e.g., from operating on behalf of the system to operating on behalf of the user).

The System Architecture criterion addresses the most critical aspect of trusted computing: the ability of the TCB to protect itself from untrusted processes. The C2 System Architecture criterion embodies three requirements:

1. The TCB must maintain for its own execution a domain (see section 3.3 below) that protects it from external interference and tampering.
2. Resources controlled by the TCB may be a defined subset of subjects and objects.

3. The TCB must isolate the resources to be protected so that they are subject to access control and auditing.

3.3 DOMAIN SEPARATION

As used in the TCSEC, the term *domain* refers to the set of objects that a subject is able to access. [14] Domain separation relates to the mechanisms that protect objects in the system. For address translation purposes, the domain separation mechanism might be execution rings, base address registers, or segmentation descriptors. In an AIS that copies files into memory, several domain-separation schemes can prevent data transfers from beyond the end of the file or accesses to arbitrary locations on the disk.

The requirement for TCB domain separation is based on the fact that if untrusted subjects are able to change the TCB, then any security mechanisms that TCB provides are useless! Therefore, this requirement addresses two essential attributes: non-tamperability and noncompromisibility. [37] Tampering generally refers to improper alterations; in this context, it involves changing the system in such a way that the intended behavior of the TCB itself is modified with respect to the enforcement of its security properties. This could happen, for example, if TCB code, data structures, or control parameters were modified. The domain of the TCB also must be self-protecting so that processes in the user domain cannot tamper with TCB code, data structures, control parameters, hardware, or firmware.

Compromise can be examined from three perspectives: compromise from above, compromise from within, and compromise from below. Compromise from above occurs when an unprivileged user is able to write untrusted code that exploits a vulnerability; e.g., finding an escape from a highly-restricted menu interface, installing or modifying a rule in an untrusted rule base that subverts a trusted rule base, or causing a denial of service. The compromise resulting from the execution of a Trojan horse (see section 4.2) that misuses the discretionary access control mechanism is another example of compromise from above. Compromise from within occurs when a privileged user or process misuses the allocated privileges, or when a programming error is made in the implementation of a trusted program. For example, compromise from within could result from a system administrator's accidentally or intentionally configuring the access tables incorrectly. Compromise from below occurs as a result of malicious or accidental failure of an underlying component that is trusted and can result from faults in the compiler or modifications to the hardware. [37]

Although the TCSEC criterion requires only that the TCB "maintain a domain for its own execution," compromise from within must be considered even for the single-layered TCB. To enable a TCB to enforce the security policy, some subjects internal

to the TCB must be "trusted;" i.e., they must run with privileges that allow them to bypass one or more of the security mechanisms. For example, the login program must run with privilege, since until it completes its function, the user on whose behalf it is running is not yet known (or at least has not been authenticated). Trusted programs must be analyzed and tested just as thoroughly as the mechanisms that enforce the security policy, to ensure that they behave as specified and do not compromise the integrity of the TCB from within.¹

An important aspect of domain separation within the CPU is "execution state" or "mode of operations." Most multi-user computer systems have at least two execution states or modes of operation: privileged and unprivileged. The TCSEC requires that the TCB maintain for itself a distinct execution state that protects it from the actions of untrusted users. Some common privileged domains are those referred to as "executive," "master," "system," "kernel," or "supervisor" modes; unprivileged domains are sometimes called "user," "application," or "problem" states. In a two-state machine, processes running in a privileged domain may execute any machine instruction and access any location in memory. Processes running in the unprivileged domain are prevented from executing certain machine instructions and accessing certain areas of memory.

Probably the most straightforward approach for implementing domain separation is to design a TCB that takes advantage of multi-state hardware; i.e., a CPU that provides two or more hardware states (rings, modes, domains). IBM's Multiple Virtual Storage/System Product (MVS/SP), Digital Equipment Corporation's VAX/VMS, and Data General Corporation's AOS/VS illustrate the diversity in hardware-based domain separation. MVS/SP provides two execution states: problem state for user programs and supervisor state for system programs. [21] VAX/VMS provides four processor access modes, which are used to provide read/write protection between user software and system software. [18] The MV/ECLIPSE architecture of AOS/VS provides eight execution "rings," ranging from ring 0 (most privileged) to ring 7 (least privileged), with the AOS/VS kernel running in ring 0 and user programs in ring 7, and with firmware-implemented gates protecting ring boundaries. [17]

For most hardware platforms, the domain separation requirement will mean that at least two hardware states are provided, where one state permits access of privileged instructions necessary to manipulate memory-mapping registers. Memory mapping alone is not sufficient to meet this requirement, but may be used to enhance hardware isolation. For example, Unisys' OS 1100 Security Release I provides domain isolation through the use of hardware and software mechanisms that include per-process virtual address spaces, per-process stacks, and hardware-based state changes. [27]

However, the multi-state mechanism need not be totally implemented in hardware.

¹Note that a "trusted process" is trusted to behave correctly only with respect to the privilege(s) it requires, and not in the general sense.

The Unisys A Series MCP/AS with InfoGuard successfully achieved a C2 rating by implementing the two-state concept with a combination of "capability-like" hardware mechanisms and TCB software, including the compilers. [26] In capability-based systems, the TCB can be protected by having TCB and user domains created when the system is initialized. Since part of the domain definition is the ability to access and modify the data structures needed for domain transition, multiple states can be created on single-state hardware.

Another approach for meeting this requirement is to have all user actions interpreted by the TCB before it acts upon them. Obviously, this entails assuring that no means exist for an untrusted user to modify the TCB. To protect against compromise from below, the requirement for domain separation implies physical protection of the hardware (even though the example cited in the TCSEC requirement is software oriented). [9]

3.4 DEFINED SUBSET

The writers of the TCSEC intended the second sentence of the System Architecture requirement to be a "grandfather clause" to enable systems designed before the TCSEC existed and add-on packages such as RACF [23] and ACF2 [15] to meet the C2 criterion even though they were not capable of controlling *all* subjects and objects in the system.

The evaluation community has interpreted this requirement to mean that:

1. Only TCB-controlled subjects can access all objects.
2. Subjects not under TCB control can access only objects that are not under TCB control.

These constraints prevent uncontrolled subjects from performing raw input-output (I/O) to (controlled and uncontrolled) devices and from accessing (controlled and uncontrolled) memory. If uncontrolled subjects were allowed to perform such operations, the TCB would be unable to enforce the system security policy with respect to controlled resources. [9]

3.5 RESOURCE ISOLATION

The third sentence of the System Architecture requirement relates to subject and object subsetting discussed in section 3.4 and simply assures that the TCB imposes

its discretionary access controls and auditing on all of the subjects and objects under its control.

Chapter 4

PROTECTION MECHANISMS

The requirements for controlled access protection comprise both mechanisms and assurances. The mechanisms are functional features designed to enforce the security policy and accountability objectives discussed in Chapter 2 and include: identification and authentication, discretionary access control, object reuse, and audit (see Table 2.1 on page 11).

4.1 IDENTIFICATION & AUTHENTICATION

Controlled access protection mechanisms ultimately are tied to the trustworthiness of the AIS's identification and authentication mechanisms. One must be able to trust the system's ability to accurately, consistently, and positively identify each user, and to maintain that positive identification throughout the user's login session. Otherwise, controlled access protection cannot be assured, and any audit data collected are rendered useless. For this reason, if the system lacks acceptable identification and authentication mechanisms, it cannot be recommended for accreditation.¹

The Identification and Authentication criterion is shown in Figure 4.1. *A Guide to Understanding Identification and Authentication in Trusted Systems* [6] discusses the identification and authentication (I&A) requirement at length and provides guidance on how to design and implement effective I&A mechanisms.

Controlled access protection seeks to control users' access to information in the AIS; specifically, information contained in objects to which users can refer by name. All forms of access control (discretionary and mandatory) rely on the system's ability to identify users and to "prove" their identity when they log onto the system, and to maintain a positive association between each individual user and the actions for

¹See Reference [38], D.7, for exception conditions.

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanisms (e.g., passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

Figure 4.1: TCSEC C2 Identification and Authentication Criterion

which he or she is responsible.

Identification is generally implemented by simply asking for a login name, usually associated in some way with the person's identity. The system checks this name against its list of authorized users. Then, to protect against an unauthorized user's masquerading as the authorized user, the system asks for some "proof" (authentication) that the user is whom he or she claims to be. Authentication generally involves one or more of three types of "proof:" (1) something the user knows (e.g., a password), (2) something the user has (e.g., an authentication device), or (3) something the user is (e.g., a retinal scan).

Most EPL products implement I&A using the simple login name and password, and this approach is acceptable. Some products strengthen their password mechanisms by enforcing rules such as aging and length requirements (e.g., Hewlett Packard's MPE V/E [19]) or case restrictions and requirements for special characters (e.g., IBM's MVS/XA with RACF [22]), or by providing random-password generators (e.g., AT&T's System V/MLS and Wang's SVS/OS [16] [28]). However, as with any mechanism, the integrity of password protection is only as strong as the integrity and responsibility of its users. Regardless of whether an AIS is built on an EPL product, the Trusted Facilities Manual (see section 5.4), the Security Features Users Guide (see section 5.5), the system administrator, and user training should all stress users' responsibilities in ensuring that their passwords are difficult to guess, protected, and changed regularly. The *Department of Defense Password Management Guideline* [13] discusses issues relating to the use of passwords for user authentication, and the *Information System Security Officer Guideline* [33] discusses user training and password management.

NSA has examined a number of subsystems designed to provide I&A, including password devices, challenge-response personal authentication devices, and biometric devices. The *Information Systems Security Products and Services Catalogue* [34] contains information regarding these devices. These products may offer an interim solution for a system that is not built on an EPL product and that lacks I&A mech-

anisms. However, the use of one or more separately-rated subsystems such as these does not imply an overall product rating as defined in the TCSEC.² Mechanisms, interfaces, and the extent of required supporting functions for each subsystem may differ substantially and may introduce significant vulnerabilities that are not present in products whose security features are designed with full knowledge of interfaces, and hardware and software support. Therefore, incorporation of one or more evaluated subsystems into an AIS is not equivalent to building an AIS on an EPL product.

4.2 DISCRETIONARY ACCESS CONTROL

Controlled access protection enforces a security policy known as *discretionary access control* (DAC), which is a means of restricting access to named objects based upon the identity of subjects and/or groups to which they belong. Systems that provide DAC assure that access to objects that are available to users (i.e., “named” objects) are controlled at the “discretion” of the user (or group) with whom the object is associated (sometimes called the “owner” of the object). The DAC criterion is shown in Figure 4.2.

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanisms (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals or defined groups of individuals, or by both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

Figure 4.2: TCSEC C2 Discretionary Access Control Criterion

Five basic mechanisms have been used to implement DAC.³

1. *Access Control Lists (ACLs)* implement an access control matrix (wherein the columns represent users, the rows protected objects, and each cell indicates the type of access to be granted for the subject/object pair) by representing the columns as lists of users attached to the protected object.

²Further, augmenting or replacing an evaluated product's I&A mechanism with a subsystem invalidates the rating.

³Some AISs may use more than one DAC mechanism; however, more is not necessarily better.

2. *Protection Bits* use a bit vector, with each bit representing a type of access. The most common example is the Unix[®] implementation of a nine-bit vector representing read, write, and execute accesses to be granted to the object's owner, its group, and everyone else.
3. *Capabilities* allow access to a protected object if the requester possesses the appropriate protected "capability," which both identifies the object and specifies the access rights to be allowed to the user who possesses that capability.
4. *Profiles* associate with each user a list of protected objects that the user may access.
5. *Passwords* associate one (all types of access) or more (different types of access) passwords with each object.⁴

A Guide to Understanding Discretionary Access Control in Trusted Systems [5] describes in greater depth each of these mechanisms and discusses issues involved in designing, implementing, and evaluating them. Most of the products evaluated to date, including Honeywell's Multics [20], DEC's VAX/VMS [18], Hewlett Packard's MPE/VE [19], Data General's AOS/VS [17], Unisys' OS 1100 [27], and IBM's MVS/SP [21], have implemented DAC through the use of ACLs. AT&T's System V/MLS [16] uses the traditional Unix[®] protection bits, and Trusted Information Systems' Trusted XENIX [25] implements both protection bits (by default) and ACLs (at the user's discretion).

DAC provides to individual users and groups the capability to specify for each of their objects (e.g., files and directories) the kinds of access the system will grant to other users and groups. This capability is very useful for both ordinary users and system administrators. It allows each user to decide for himself or herself what individuals and groups of individuals the system should allow to read, write, or execute the directories and files he or she creates. System administrators commonly use DAC to protect system directories and files so that ordinary users can read or execute (or search, in the case of directories) them, but only system administrators can modify them. For example, DAC enables ordinary users to spool print jobs (i.e., write into the print queue) but does not allow them to read, reorder, modify, or remove other users' queued jobs. Only a program acting on behalf of a user or group with system privileges (i.e., individual or group to which the print queue belongs) can perform these actions.

However, most DAC implementations contain a flaw that renders them susceptible to *Trojan horses*. This is due to the fact that when a user executes a program, it runs with the DAC accesses of that user. This enables the following scenario to occur.

[®]Unix is a trademark of Unix System Laboratories, Inc.

⁴Passwords generally are not considered an acceptable implementation of DAC.

1. Dan Devious writes a program that performs a very useful function, say travel expense accounting, and attaches some lines of code that copy all of the files in the *mail* directory of the user who executes it into a directory that Dan owns.
2. Dan gives everyone execute access to his program and tells everyone about its utility. (He also gives everyone write access to his directory, but does not mention this.)
3. Nick Naive executes Dan's program to calculate his travel expenses. The program works just as Dan described it, and Nick is elated. However, unknown to him, the program has also copied all of Nick's mail files into Dan's directory!

Because of this vulnerability and the "discretionary" nature of DAC, this access control mechanism is not useful for segregating objects with different classification levels or categories. Mandatory access control mechanisms are necessary to provide classification-level separation.

Some operational systems have attempted to use DAC to enforce strict need-to-know separation by assigning different need-to-know categories to different groups. DAC is neither intended to be, nor effective as, a mechanism for strictly enforcing need-to-know separation. Under DAC, any user who has or can usurp the appropriate permission is able to transfer access rights to another user to whom direct access would otherwise be forbidden. The following two examples illustrate how this might occur.

1. George puts the results of his latest project experiment into *georges_data*. To ensure that Zelda and Fran, who are working on the same project and assigned to group **project**, can read the results, he assigns it the ACL shown in Figure 4.3.

project	read
others	no access

Figure 4.3: ACL for File *georges_data*

Zelda wants to share George's results with her friend Neil, who is not working on the project. So she copies *georges_data* into a file named *zeldas_data* and sets its ACL to allow both herself and Neil to read it. She then tells Neil where he can find the file, and he continues to spread access to others in a similar manner.

While this ACL may look like it would provide the needed protection, "read" access also enables any user in group **project** to copy *georges_data* into another

file with its own ACL and to assign to it whatever accesses that user wishes. Thus a file whose contents are intended to be protected from disclosure can be disclosed to supposedly “unauthorized” users.

2. On most Unix[®] systems, typing “ls -lga” (list all entries in long format, giving mode, number of links, owner, group, size in bytes, and time of last modification) in directory *study* produces the output shown in Figure 4.4.

drwxrwx---	2	sally	hackers	512	Aug 22 20:44	./
drwx--x---	4	sally	users	3584	Apr 24 11:57	../
-rw-r-----	2	sally	hackers	514	Sep 19 13:33	progress

Figure 4.4: Output from Directory Study

Group **hackers** includes Ted, Sally, and Ollie. Ted wants to modify Sally’s *progress* file, but she has given him (i.e., group **hackers**) only read permission. Although Ted does not have write access to *progress*, he knows that since he has write access to its containing directory *study* and read access to the file, he can give himself write access by executing the sequence of commands shown in Figure 4.5 to virtually change the file’s permission bits.

cat progress > newprogress	#Copy the contents of file progress to
	#file newprogress
rm progress	#Remove file progress
mv newprogress progress	#Rename ‘‘newprogress’’ ‘‘progress’’
chmod 660 progress	#Change accesses to progress to allow
	#owner and group to read and write it

Figure 4.5: Unix Command Sequence

In this case, Sally believes she has sufficiently protected her file *progress* so that only she is able to write to it. However, because group **hackers** has read access to the containing directory, any user in group **hackers** is able to see that a file named *progress* exists. Further, write access to directory *study* enables any user of group **hackers** to modify the directory’s contents. So any user in group **hackers** is able to add files to and delete files from *study* and to virtually change the DAC permission on any of its files to which they have read (i.e., copy) access. Thus, any user in group **hackers** can modify Sally’s *progress* file.

As is apparent, reliance on DAC control could very quickly result in a breakdown of need-to-know protection. While an AIS with mandatory access controls could contain the same DAC vulnerability, those controls would confine the propagation to a single classification level and category. DAC should *not* be used for separation that requires strong enforcement and assurance.

4.3 OBJECT REUSE

One could view the Object Reuse criterion shown in Figure 4.6 as a “negative” requirement in that it requires that something be “*not* present.” To meet the object reuse criterion, the AIS must ensure that no information generated by one user’s process is available to the next user’s process when the object containing that information is reallocated.

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB’s pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject’s actions is to be available to any subject that obtains access to an object that has been released back to the system.
--

Figure 4.6: TCSEC C2 Object Reuse Criterion

Note that the object reuse criterion refers to “storage” objects, as contrasted with the “named objects” to which the DAC criterion applies. A *storage object* is an object that supports both read and write accesses and may or may not be “named.” A *Guide to Understanding Object Reuse in Trusted Systems* [7] explains the object reuse criterion and provides guidance on how to design and incorporate effective object reuse mechanisms into an AIS.

The objective behind the object reuse requirement is to prevent information from being inadvertently (and by extension, deliberately) disclosed to users not authorized to see it. In contrast with the DAC mechanism, which seeks to protect the containers of information (i.e., named objects), the object reuse requirement seeks to protect the information contained in the AIS’s storage objects. Thus object reuse requires that each container be initialized before it is allocated to a subject.

However, although the level of abstraction at which the object reuse mechanism is implemented is that of storage objects, ensuring complete and effective implementation requires consideration of how named objects are mapped into physical storage objects. The object reuse guideline describes a methodology for doing this.

A number of approaches for meeting the object reuse requirement exist and are specific to the storage objects being considered. Whether the object reuse mechanism operates

at allocation or deallocation is left to the discretion of the implementer. The system may initialize a storage object any time between when it releases the object when it reallocates it. However, if the system does not initialize the object immediately, it must protect as a system resource any information it contains. Table 4.1 identifies some examples of possible object reuse mechanisms. Note that a given type of storage object may require one or more mechanisms. The object reuse guideline discusses these mechanisms more fully.

Storage Object	Implementation
Primary Storage (e.g., random access memory, cache, translation buffer)	<ul style="list-style-type: none"> • Overwriting memory page with fixed or random pattern and/or (for efficiency) new data
Fixed Media (e.g., fixed disk, terminal, operator console)	<ul style="list-style-type: none"> • Overwriting physical data blocks • Purging associated entries in page management table • Purging directory information residing on media
Removable Media	<ul style="list-style-type: none"> • On-line overwriting with approved fixed or random pattern • Degaussing^a • Off-line overwriting

Table 4.1: Object Reuse Mechanisms

^aFor further information regarding data remanence products, see *A Guide to Understanding Data Remanence in Automated Information Systems*. [3]

4.4 AUDIT

The Audit criterion requires the capability to collect information regarding system events, thus supporting the monitoring of system use and the investigation of possible attempts to breach security. Importantly, the Audit criterion, shown in Figure 4.7

on page 30 requires that the AIS be *capable of* auditing, and not that the system actually *perform* auditing. The accreditor is responsible for determining what events the system must audit and any additional mission-specific audit requirements. The Information System Security Officer (ISSO) or designated auditor is responsible for configuring and administering audit.⁵

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of access to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into the user's address space (e.g., file open, program initiation), deletion of objects, actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. For each recorded event, the audit record shall identify: data and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity.

Figure 4.7: TCSEC C2 Audit Criterion

Audit features provide the capability to record, examine, and review security-relevant activities on the system either as they are occurring or retrospectively. The capability to perform real-time auditing is not among the minimal requirements for controlled access protection.⁶ Rather, the system must provide the capability to configure the system to audit the set of events the ISSO specifies, to present this information in a manner that is useful in investigating security incidents after they have occurred, and to monitor users' actions in order to anticipate and potentially neutralize impending security attacks.

A Guide to Understanding Audit in Trusted Systems [1] discusses five objectives of the audit mechanism:

1. To allow review of patterns of access to individual objects, access histories of specific processes and users, and the use of various protection mechanisms and their effectiveness.

⁵The *Information System Security Officer Guideline* [33] provides guidance to ISSOs in configuring audit mechanisms to audit the required events, and in reviewing and maintaining audit trails.

⁶However, some products, such as DEC's VAX/VMS [18], do provide some real-time monitoring/alarming capability.

2. To detect repeated attempts to bypass protection mechanisms.
3. To monitor use of privileges.
4. To deter habitual attempts to bypass the system protection mechanisms (which requires that users know that their actions are being audited).
5. To provide additional assurance that the protection mechanisms are working.

As pointed out in section 4.1, the integrity of the audit mechanism is highly dependent upon the integrity of the I&A mechanisms. Unless the system positively identifies users, it cannot correctly associate their actions with them, and no audit mechanism can be effective. As with all controlled access protection mechanisms, the TCB must implement the audit-collection function, and only ISSOs or their designees should be able to enable or disable auditing, and to configure the audit mechanism (i.e., to set the events to be recorded, the users for which data are to be collected, etc.) in accordance with the security policy. The TCB must protect the data the audit mechanism collects; only audit personnel should be able to read audit data. Further, the TCB must protect the audit trail from unauthorized modification and from loss due to overwriting (such as might occur if a circular file were used to store audit data), exhaustion of physical memory reserved for storage of audit data, or a system crash.

The system must be able to record the following types of events:

- Use of identification and authentication mechanisms (i.e., login).
- Introduction of objects into a user's address space (e.g., file open, file creation, program execution, file copy).
- Deletion of objects from a user's address space (e.g., file close, completion of program execution, file deletion).
- Actions taken by computer operators and system administrators and/or system security administrators (e.g., adding a user).
- All security-relevant events (e.g., use of privileges, changes to DAC parameters).
- Production of printed output.

For each auditable event, the TCB must be able to record the following information:

- Date and time of the event.
- Unique identifier of the user on whose behalf the subject generating the event was operating.

- Type of event (one of the above).
- Success or failure of the event.
- Origin of the request (e.g., terminal identifier) for identification and authentication events.
- Name of the object that was introduced into or deleted from the user's address space.
- Description of actions taken by the system administrator (e.g., modifications to the security databases).

The ISSO or designee must be able to audit based on individual identity and on object identity. Whether the system allows the ISSO to pre-specify individuals and/or objects, or provides a post-processor to extract data associated with specified individuals and/or objects, is a design decision. From a security perspective, either approach could be deemed acceptable.⁷ Data compression and reduction tools are also desirable (but not required) features. A number of vendors have implemented extensive audit-processing capabilities in their products. For example, Prime Computer, Inc.'s Primos [24] and Unisys Corporation's OS 1100 Security Release I [27] provide auditing facilities which include collection, reduction/reporting, backup, and crash-recovery capabilities.

⁷Note, however, that the post-processing option may result in an audit-collection mechanism that overly burdens the system, resulting in a tendency to turn auditing off entirely.

Chapter 5

DOCUMENTATION AND LIFE-CYCLE ASSURANCE

A number of requirements are derived not from the security policy *per se*, but from the assurance control objective (see Table 2.1 on page 11) and from the needs for evaluation evidence and documentation to support continuing maintenance of the evaluated trust. This chapter discusses these documentation and life-cycle support requirements.

5.1 DESIGN DOCUMENTATION

The Design Documentation criterion, shown in Figure 5.1, focuses on the need to document coverage of the protection philosophy. While this information is useful in understanding how the system provides trust, it is not sufficient to enable an analyst to understand the design of the AIS. More detailed design documentation is needed to ensure that the system can be understood and maintained securely.

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. If the TCB is composed of distinct modules, the interfaces between these modules shall be described.
--

Figure 5.1: TCSEC C2 Design Documentation Criterion

The primary purposes of design documentation are:

- To help evaluators (e.g., NSA product evaluators, technical analysts) achieve a sufficient understanding of the system to enable them to assess the complete-

ness and correctness of the design, and to give them enough confidence in the developer's understanding and capabilities to warrant a recommendation that the system be approved (e.g., for an NSA rating or DAA accreditation).

- To enable developers and maintainers to understand the design of the AIS well enough so that they can make any necessary changes to the AIS without adversely affecting the system's trustworthiness.

In order to serve these purposes, the design documentation must describe all of the protection mechanisms of the TCB. In other words, the design documentation must accurately and completely describe all of the software, firmware, and hardware components and how they work together. These descriptions should be in sufficient detail to enable an evaluator, system programmer, or certifier to understand the security design and implementation such that he or she can predict the security impacts of a hypothesized or proposed modification.

As discussed in Chapter 3, each conceptual "layer" of the TCB must be trustworthy from the perspective of its overlying layers. The hardware and software design documentation needs to clearly describe how this trustworthiness is assured. For example, the hardware design documentation should describe the interface between the hardware and the operating system in sufficient detail to enable someone analyzing the system to feel assured that the TCB cannot be circumvented (i.e., compromised from below), enabling an unprivileged user to gain direct access to the system's physical resources (e.g., disk blocks, physical I/O). Similarly, the software design documentation must describe how the TCB provides self-protection and isolation from user processes (i.e., prevents compromise from within and from above).

Good design documentation describes how the protection mechanisms relate to the overall architecture of the system. *A Guide to Understanding Design Documentation in Trusted Systems* [4] provides guidance that developers can use in assuring that their design documentation is acceptable, and that analysts can use in their evaluation.

5.2 SYSTEM INTEGRITY

The System Integrity criterion, shown in Figure 5.2, is levied upon the hardware and firmware components of the TCB.

"Integrity" implies that something is maintained in an unimpaired condition, and *system integrity* implies that an AIS and the system data upon which its operation depends are maintained in a sufficiently correct and consistent condition. [37] The intent of the system integrity requirement is to ensure that some mechanism exists to validate the correct operation of all TCB hardware and firmware (including peripheral devices).

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

Figure 5.2: TCSEC C2 System Integrity Criterion

Typically, the first time this requirement comes into play is at system boot time. The system should provide some mechanism for assuring that the TCB (i.e., all security-relevant hardware and firmware, including peripheral devices) is initialized correctly. This should not impose a problem for most systems, since most commercially available computer systems provide a mechanism and procedures for performing a comprehensive diagnostic routine when they are powered on.

The system also should provide mechanisms for periodically validating the correct operation of its hardware and firmware. For example, tools for performing comprehensive diagnostics following preventive maintenance actions and to ensure secure system shut-down should be available. Documentation describing the functionality and operations of all integrity mechanisms should be provided.

5.3 CONFIGURATION MANAGEMENT

Changes to an existing AIS are inevitable, and the purpose of configuration management (CM) is to ensure that these changes take place in a controlled environment and that they do not adversely affect any trust properties of the system. CM provides assurance that additions, deletions, and changes to the AIS do not compromise its inherent trust. CM therefore is of critical importance with regard to life-cycle assurance. During development and in operation, the AIS's software and hardware must not be changed improperly or without authorization, control, and accountability.

The TCSEC does not specify a Configuration Management criterion for classes lower than B2. However, the AIS organization should recognize the important role that CM plays both in performing the technical analysis and in assuring the continued secure operation of the system. Although CM is not a controlled-access-protection requirement, requiring sound CM policy and procedures, and subjecting them to technical assessment, are strongly recommended.

AISs being analyzed for certification and accreditation should provide documentation and compliance evidence demonstrating that an effective CM program exists and that configuration control is enforced.

A Guide to Understanding Configuration Management in Trusted Systems [2] discusses the Configuration Management criterion imposed on products submitted for

a B2 or above rating and provides a good overview of the CM process and the functions involved: configuration identification, configuration control, configuration status accounting, and configuration audit. MIL-STD-483, *Configuration Management Practices for Systems, Equipment, Munitions, and Computer Programs* [12], provides CM standards to be applied to DoD systems.

Suggested items to cover in the AIS's CM plan are:

- Unified discussion of configuration control as implemented by the developer; description of the process for handling a change from entry into the process through final approval and implementation.
 - Description of the approach used to determine configuration items (CIs), including a rationale for the chosen granularity.
 - Naming conventions for CIs.
 - Policies for creating new CIs or changing CIs.
 - Decomposition of the following system components into CIs, with unique identifiers for each:
 1. The TCB.
 2. Any hardware and/or software features that are used to periodically validate the correct operation of the TCB.
 3. The Security Features User's Guide.
 4. The Trusted Facility Manual.
 5. The test plan, the test procedures that show how the security mechanisms were tested, and the expected results of the security mechanisms' functional testing.
 6. The design documentation.
 7. The CM Plan.
- Explanation of the results of the preliminary screening of proposed changes and a discussion of any identified potential effects on the TCB.
- Description of safeguards against the incorrect categorization of changes.
- Detailed discussion of security analysis for changes affecting the TCB.
- Description of how the Configuration Control Board (CCB) coordinates security and design analyses and reviews system changes, including CCB composition, lines of authority, and identification of security specialists and their roles.
- Description of the content of engineering change orders and a discussion of how they are generated and handled within the CM system.

- Description of procedures for assuring that all approved changes are implemented correctly and that only approved changes are made, including the structure and interactions of the implementation and test groups and the management of system code.
- Description of the nature and operation of the Configuration Review Board (CRB).
- Discussion of the final review process.
- Identification of any limitations or constraints on the CM process.

5.4 TRUSTED FACILITY MANUAL

No matter how strong the security architecture and mechanisms are, and how trustworthy the users are, an AIS's "weakest link" is its administration and operations. Even if the AIS is built on an EPL product, the protection the product is capable of delivering is actually provided only if the system is configured in one of the evaluated configurations indicated in the product's EPL entry and is operated as described in the Trusted Facility Manual (TFM). The TFM criterion shown in Figure 5.3 addresses this critical need.

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given.

Figure 5.3: TCSEC C2 Trusted Facility Manual Criterion

The TFM is written for AIS administrators (e.g., ISSOs) responsible for configuring, operating, and monitoring the system and for investigating potential violations of the security policy. For some systems (in particular, products rated B3 and A1), the administrative role is broken down into unique privilege classes (e.g., operator, security administrator, auditor). However, for controlled access protection, a single privileged role is acceptable. This fact renders the TFM even more important.

Guidelines for Writing Trusted Facility Manuals [32] provides a detailed discussion of the TFM criterion and the important role the TFM plays in ensuring the trustworthiness of the system, and *Information System Security Officer Guideline* [33] discusses the overall role of the ISSO. The TFM generally is not intended to be part of the DAA accreditation package, but is required for controlled access protection and should be examined during the technical analysis. The TFM is prepared to support life-cycle

trusted system operations, and its goal is to provide detailed, accurate information on how to:

1. Configure and install the system to a secure state.
2. Operate the system in a secure manner.
3. Make effective use of the system privileges and protection mechanisms to control access to administrative functions and databases.
4. Avoid pitfalls and improper use of administrative functions that would compromise the TCB and user security.

TFMs distributed with EPL products contain information addressing these goals, and if the AIS is built on an EPL product, this document should be part of the system's TFM. In addition, the system's TFM should contain information regarding site-specific operations, including the security policy to be enforced in configuring and operating the AIS in its unique environment under both routine and emergency situations.

5.5 SECURITY FEATURES USER'S GUIDE

Whereas the TFM is written for system administrators, the *Security Features Users Guide* (SFUG) is written for the general, unprivileged users of the AIS. The SFUG criterion is shown in Figure 5.4. Using terminology a user unfamiliar with the operating system can understand, the SFUG should describe the security mechanisms the system provides to the general user. For example, the SFUG should explain how login works, provide guidance and warnings regarding the selection and use of passwords, explain how to set the DAC permissions on files and directories, and briefly discuss the role auditing plays in the operation of the AIS. The objective of the SFUG is to provide information and warnings to help assure that the system's protective features are used appropriately and consistently.

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, guidelines on their use, and how they interact with one another.

Figure 5.4: TCSEC C2 Security Features User's Guide Criterion

A Guide to Writing the Security Features User's Guide for Trusted Systems [8] provides guidance for potential authors of SFUGs and includes some illustrative anno-

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. Testing shall be done to assure that there are no obvious ways for an unauthorized user to bypass or otherwise defeat the security protection mechanisms of the TCB. Testing shall also include a search for obvious flaws that would allow violation of resource isolation, or that would permit unauthorized access to the audit or authentication data.

Figure 5.5: TCSEC C2 System Testing Criterion

tated outlines.¹

5.6 TESTING

The final step in the technical analysis (see Chapter 6) is testing, which includes both test planning and running the functional tests. The test objective with respect to controlled access protection is to ascertain whether the documented security mechanisms work as they are described. Note that the TCSEC System Testing criterion (see Figure 5.5) requires assurances that no “obvious ways” exist to bypass or otherwise defeat the security protection mechanisms, and a search for “obvious” flaws. Thus, the technical analysis to support certification involves testing to ensure that the documented security functionality exists and works as claimed; this level of testing does not require an in-depth penetration effort, which would involve the generation of hypotheses to ascertain whether “non-obvious” penetrations are possible.

Note that the TCSEC does not precisely define “obvious,” and what is “obvious” to one analyst may be enigmatic to another. The analysts should interpret “obvious” based on the identified threats to the system. For example, some Unix[®] vulnerabilities that are well-known (i.e., “obvious”) within campus computing centers may be far less threatening (i.e., “obvious”) in a closed DoD environment.

The analysts should conduct functional testing at the user interface of the system. That is, they should test all of the security functionality available to the general, unprivileged user. All of the mechanisms discussed in Chapter 4 should be tested to ensure that they do what they are intended to do and that they do not contain “obvious” flaws in their design or implementation. If the system is built on an EPL product, the test suite provided with the product may be useful for this purpose. Further, the system integrity mechanisms discussed in section 5.2 should be tested to ensure that they work as claimed.

¹User training is as important as user documentation. *Information System Security Officer Guideline* [33] provides some guidelines for user training.

Chapter 6

TECHNICAL ANALYSIS

6.1 SELECTION OF ANALYSTS

A team of qualified individuals should be selected to analyze the AIS to ensure that it provides the required levels of controlled access protection. All members of the team should have the equivalent of at least a bachelor's degree in Computer Science or Computer Engineering. At least one team member should possess technical expertise in computer hardware architectures, and all members should possess technical expertise in operating systems. All team members should be familiar with and understand security issues related to computer hardware and operating systems. In addition, the analysts should understand the system's mission, its environment, its security policy, and its identified threats.

Before beginning the technical analysis, all members of the team should have received training in the methodology described in this document and in the operations and internal architecture of the AIS to be analyzed. If the system is built on an EPL product, the analysts should have obtained and become familiar with the product's Final Evaluation Report.¹ All team members should feel comfortable on the system as both administrators and general users and should be able to design and implement test programs for the system.

¹The product's EPL entry will contain the title and document number of this report, which can be requested from the NTIS.

6.2 TECHNICAL ANALYSIS PROCESS

Figure 6.1 depicts the steps (described below) involved in performing a technical analysis of an AIS to ensure that it provides the functionality and assurances necessary for controlled access protection. Although this process is correct and complete with respect to its objectives, it cannot predict and does not address many issues that may arise when analyzing a complex system (e.g., issues relating to the composition of networks). Also note that the order of some steps of the process are arbitrary and could be conducted in a different order or in parallel (e.g., DAC and audit assessments). Steps in which dependencies exist and order is important are identified. As noted above, the analysts should have a clear understanding of the system's mission and policy, security requirements, concept of operations, and operational environment before beginning this process.

In the process flow shown in Figure 6.1, each rectangle represents an activity, and each edge represents a possible course of action, with the conditions associated with that action noted alongside the edge. For every activity, only one set of entry and exit conditions applies in any given instance. If an incoming conditional arc (i.e., one on the left side of a rectangle) is labeled "OR," then the occurrence of one of the edges associated with that conditional will result in the activity's being initiated. If an outgoing conditional arc (i.e., one on the right side of a rectangle) is labeled "OR," then the activity effects one of the actions identified on the outgoing edges.²

Each "Fix" task is assumed to include the CM process, which will assure that the correction does not adversely affect preceding analyses. If a fix affects a mechanism that has already been analyzed, the process should revert to the point at which the affected mechanism is analyzed. For example, if a fix to correct an audit deficiency affects the implementation of I&A, the analysis should return to the "Assess I&A" task.

The *Trusted Product Evaluation Questionnaire* [40] is referenced frequently in the following task descriptions. This questionnaire was designed as an instrument for gathering from vendors preliminary information about products submitted to NSA for evaluation. However, the referenced items are equally applicable in the context of this analysis.

As this process flow shows, by far the easiest and most direct way to attain controlled access protection is to build the system on a product that has been evaluated by NSA and rated C2 or higher (assuming it is correctly configured, including no modifications to the TCB).

²This notation also will accommodate "AND" conditions, but because none of these conditions appear in the diagram, they are not defined here.

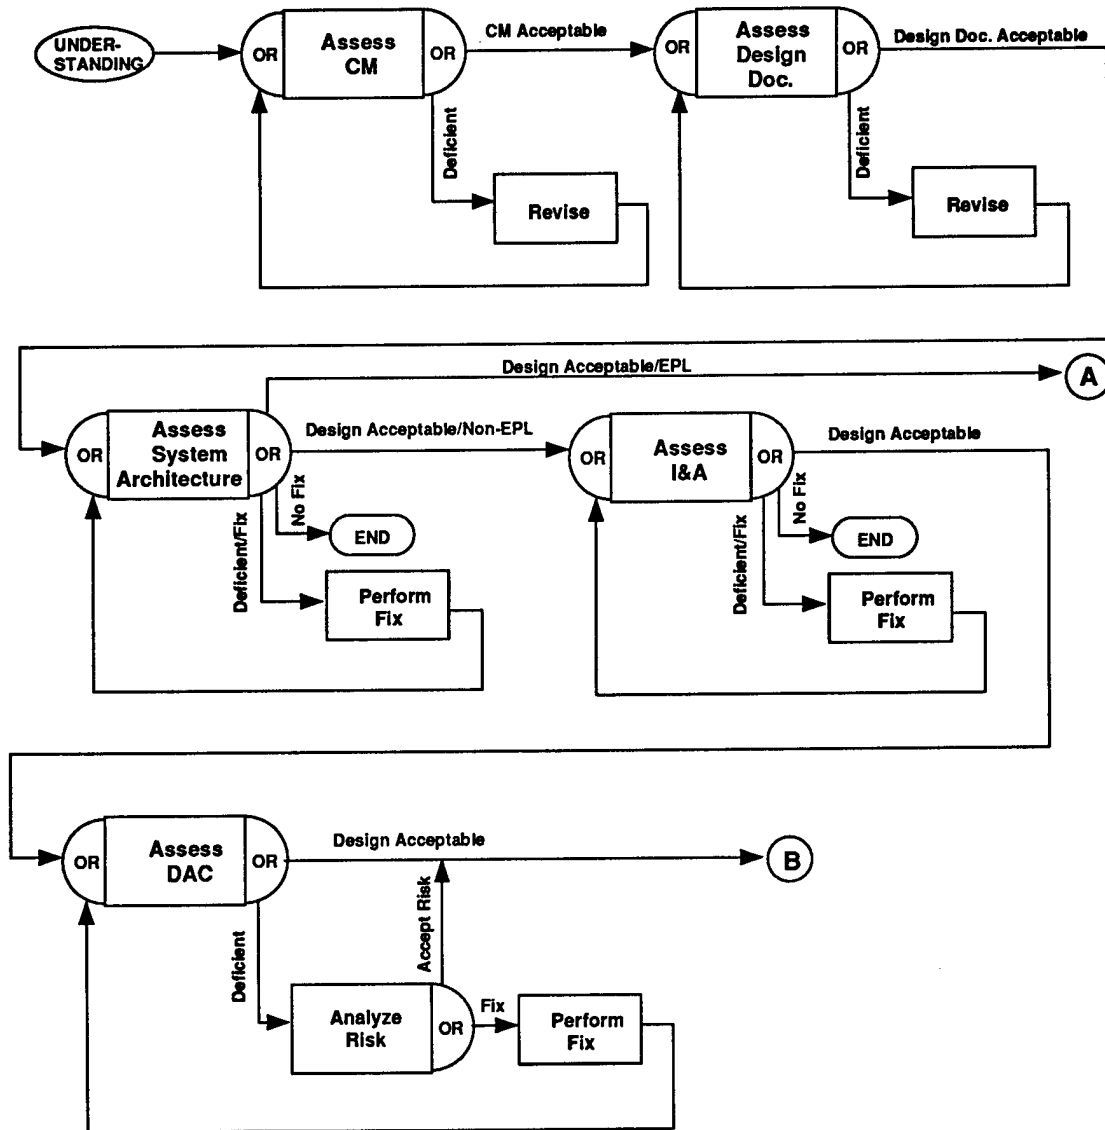


Figure 6.1: Controlled Access Protection Technical Analysis Process

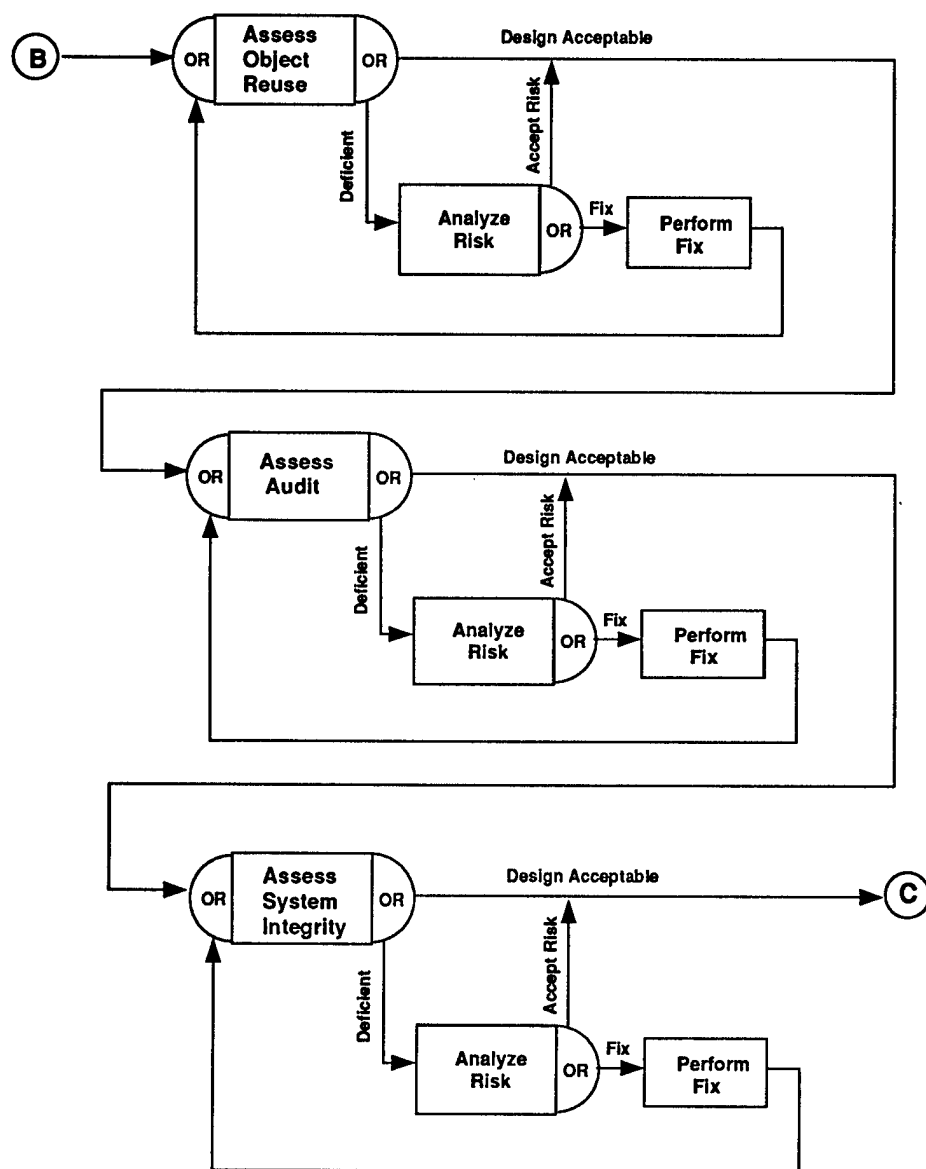


Figure 6.1: (cont.) Controlled Access Protection Technical Analysis Process

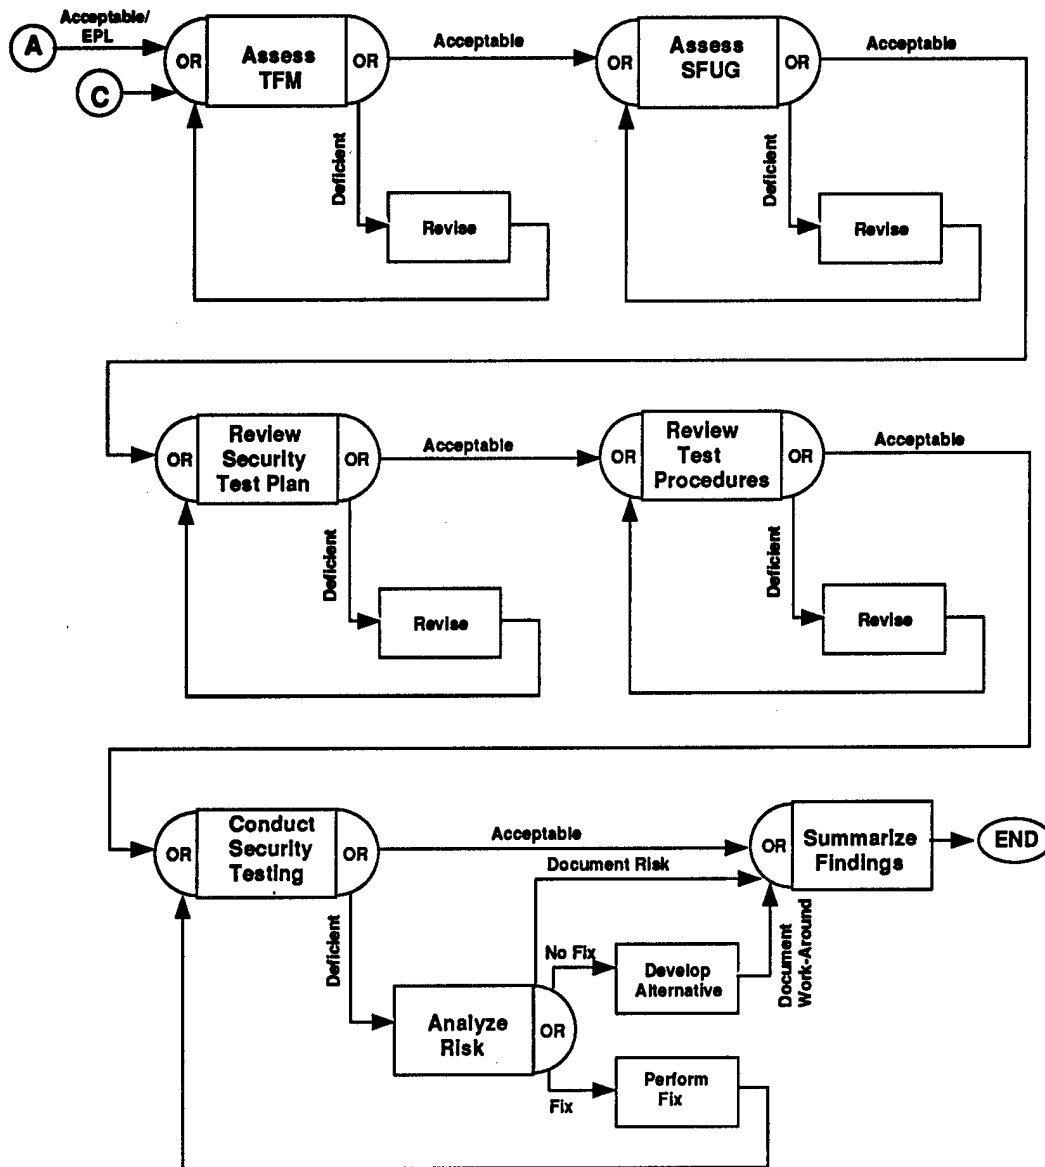


Figure 6.1: (cont.) Controlled Access Protection Technical Analysis Process

Step 1. Assess Configuration Management. The first step in the assessment is to gain assurance that a sound configuration management program is in place. This step should be performed before any analysis of the system itself begins to ensure that all changes that are made to the system software and documentation are controlled. The configuration management requirement is discussed in section 5.3. The analysts review the documentation describing the plans and procedures for providing CM³ and control, and complete items 1 through 4 in Section 2.13 of the *Trusted Product Evaluation Questionnaire*. An acceptable CM system will cover all of the items discussed in section 5.3.

The analysts ascertain whether the CM system as documented is acceptable and is enforced as documented; if not, the developer changes the CM program as required.

Step 2. Assess Design Documentation. The second step, which must be performed before and in parallel with the system architecture assessment, is to review the design documentation. Regardless of whether an EPL product is used, the analysts evaluate the hardware and software design documentation to gain an understanding of the system and to determine whether it meets the Design Documentation criterion shown in Figure 5.1 and discussed in section 5.1.

The analysts ensure that the design documentation for the hardware and software addresses all of the functionality needed to support the security policy enforced by the AIS. To ascertain whether this requirement is met, the analysts answer item 1 in Section 2.3, items 1 and 15 in Section 2.4, and item 1 in Section 2.14 of the *Trusted Product Evaluation Questionnaire*.

If the design documentation is incomplete or deficient, it is developed and revised until it accurately and completely describes the system's design and implementation.

Step 3. Assess System Architecture. The next step, which should be performed before and in parallel with analyzing the security control mechanisms, is to gain a thorough understanding of the system architecture. During this step, the analysts become familiar with the architectural foundation upon which the security mechanisms are built and determine whether the AIS meets the System Architecture criterion discussed in Chapter 3 and shown in Figure 3.3. If the security policy for the AIS includes more than controlled access protection, the analysts also need to determine how the extension to the security policy fits into the overall security architecture. For example, many DoD systems are designed to

³Although this analysis addresses CM relative to the TCB only, all applicable programs and documentation should be controlled within the CM system.

provide a restricted user interface comprising a set of menus from which an operator (unprivileged user) selects the function he or she wishes to perform, response fields or windows in which the operator enters requested data, and output fields or windows, where output and system status messages may appear. These restricted interfaces may be implemented by an untrusted application built on top of the TCB (i.e., without modifying the operating system) or as an extension to the TCB. The analysts must examine the implementation to determine which method is used. If the restricted interface is an unprivileged program residing in the user domain (see discussion in Chapter 3), then the analysts must ensure that its discretionary access control (see section 4.2) settings are correct and that it is included in system testing, but need make no assertions regarding its trustworthiness relative to the overall system architecture. If the interface is part of the TCB interface, then its mechanisms and assurances should be analyzed along with (and in addition to) the mechanisms and assurances discussed in this guideline.⁴

- If the system is built on a product rated C2 or above on the EPL, the analysts can assume that an NSA evaluation team has conducted an in-depth analysis of the vendor's proprietary design documentation and has determined that the product meets the System Architecture requirement. At this point, the analysts need to ensure that all of the following conditions are satisfied:
 1. The system is built on the evaluated configuration.
 2. The TCB has not been modified (i.e., no modifications to system code have been made, and no applications use privileged system calls intended only for internal TCB use). (Answer questions 5 and 6 in Section 2.13 of the *Trusted Product Evaluation Questionnaire*.)
 3. The mechanisms discussed in Chapter 4 are configured in accordance with the Trusted Facility Manual (see section 5.4) and the AIS's security policy.

If any of these conditions does not hold, and the deficiency cannot be corrected, the process proceeds as if a non-EPL product were used. If all of these conditions are satisfied, the analysis proceeds to step 6.

- If the system is *not* built on an EPL product or is built on an EPL product in other than its evaluated configuration, the analysts begin the architecture evaluation by completing the C1 and C2 items in

⁴This requires modification of the TCB, so if a C2-rated product is used, its rating is invalidated, and it must be analyzed as if an unevaluated product had been used. However, information contained in the Final Evaluation Report for the evaluated product will be useful in the evaluation process.

Sections 2.1 and 2.2 and items 5 and 6 in Section 2.13 of the *Trusted Product Evaluation Questionnaire* [40] to gain a full understanding of all of the subjects and objects in the system.

The analysts then attempt to gain a full understanding of the hardware and software upon which the system's protection mechanisms depend. The analysts review the design documentation (see section 5.1) for the hardware and software and complete items 2 through 10 in Section 2.3 and items 16 through 18 in Section 2.4 of the *Trusted Product Evaluation Questionnaire*.

As noted in section 1.3, a precondition of NSA evaluation is that the vendor must sign an MOU giving NSA evaluators access to highly proprietary hardware and software design documentation. Had the system been built on an EPL product, the analysts could have assumed that an NSA evaluation team had conducted an in-depth analysis of the vendor's proprietary design documentation and had determined that the product met this requirement. However, because the system is not built on an EPL product or is built on an EPL product in other than its evaluated configuration, the analysts may not have access to detailed proprietary design documentation. In this case, they will need to rely on commercial documentation distributed with the unevaluated product, documentation provided by the contractor, information they are able to glean from talking with the vendor and contractor, and any other available information.

This limited visibility into the hardware and software design is one of the most critical constraints associated with analyzing a system built on a non-EPL product or an EPL product in other than its evaluated configuration.

During this analysis, the analysts ascertain whether the System Architecture criterion is met (see Chapter 3) and whether the features discussed in Chapter 4 are present.⁵ If any deficiencies are noted and are judged "fixable," the developer ensures that the necessary modifications are made to both the system architecture and the design documentation, and the technical analysis is reinitiated. *If uncorrectable system architecture deficiencies are identified, the AIS is deemed unacceptable, and the technical analysis is terminated with a recommendation that the system not be certified or accredited* (see Reference [38], D.7, for exception conditions).

⁵At this point in the analysis, the analysts do not need to evaluate these mechanisms, but just check to ensure that they are present.

Step 4. Perform Fixes. If correctable architectural deficiencies are identified, the developer ensures that the necessary modifications are implemented.

Step 5. Perform Steps Required for Non-EPL-Based Systems. If the system architecture is acceptable, and the system is built on a non-EPL product or an EPL product in other than its evaluated configuration:

1. **Assess Identification and Authentication.** The analysts study the design to determine how well the AIS meets the I&A criterion shown in Figure 4.1 and discussed in section 4.1.

To determine whether the system meets this requirement, the analysts complete the C1 and C2 items in Section 2.6 of the *Trusted Product Evaluation Questionnaire*. If any deficiencies are noted, the analysts determine whether they can be corrected; if so, the required changes are implemented. If the AIS does not provide acceptable I&A mechanisms, it must be deemed unacceptable, since this mechanism is of critical importance to controlled access protection. Unless users are positively identified and authenticated, the TCB cannot assuredly enforce the security policy. Therefore, *if the AIS lacks acceptable I&A mechanisms, the technical analysis is terminated with a recommendation that the system not be certified or accredited* (see Reference [38], D.7, for exception conditions).

2. **Assess Discretionary Access Control.** The analysts study the design to determine how well the AIS meets the DAC criterion shown in Figure 4.2 and discussed in section 4.2.

To ascertain whether the system meets this requirement, the analysts complete the C1 and C2 items in Section 2.5 of the *Trusted Product Evaluation Questionnaire*. If any deficiencies are noted, the analysts assess the associated risks (see section 7.2) and take appropriate action as shown in Figure 6.1.

3. **Assess Object Reuse.** The analysts study the design to determine how well the AIS meets the Object Reuse criterion shown in Figure 4.6 and discussed in section 4.3.

To ascertain whether the system meets this requirement, the analysts complete the C2 items in Section 2.7 of the *Trusted Product Evaluation Questionnaire*. If any deficiencies are noted, the analysts assess the associated risks (see section 7.2) and take appropriate action as shown in Figure 6.1.

4. **Assess Audit.** The analysts study the design to determine how well the AIS meets the Audit criterion shown in Figure 4.7 and discussed in section 4.4.

To ascertain whether the system meets this requirement, the analysts complete the C2 items in Section 2.8 of the *Trusted Product Evaluation Questionnaire*.

If any deficiencies are noted, the analysts assess the associated risks (see section 7.2) and take appropriate action as shown in Figure 6.1.

5. **Assess System Integrity.** The analysts study the design to determine how well the AIS meets the System Integrity criterion shown in Figure 5.2 and discussed in section 5.2.

To ascertain whether the system meets this requirement, the analysts answer questions 1 through 3 in Section 2.11 and item 7 in Section 2.13 of the *Trusted Product Evaluation Questionnaire*.

If any deficiencies are noted, the analysts assess the associated risks (see section 7.2) and take appropriate action as shown in Figure 6.1.

Step 6. Evaluate Trusted Facility Manual. The analysts evaluate the documentation provided to instruct system administrators in how to configure and operate the AIS securely, as required in the TFM criterion shown in Figure 5.3 and discussed in section 5.4.

To ascertain whether this requirement is met, the analysts answer the following questions.

- If the system is built on an EPL product:
 - Is the evaluated TFM included in the TFM for the AIS?
 - Does the TFM stress the importance of configuring the system into its evaluated configuration and provide pointers to procedures for accomplishing this?
 - Does the TFM provide appropriately highlighted site-specific warnings with respect to actions that would invalidate the NSA rating?
 - Does the TFM provide site-specific procedures for collecting, reviewing, analyzing, and storing audit information?
 - Item 15 in Section 14 of the *Trusted Product Evaluation Questionnaire*.
- If the system is not built on an EPL product or is built on an EPL product in other than its evaluated configuration:
 - Items 11 through 16 in Section 2.14 of the *Trusted Product Evaluation Questionnaire*.

If any deficiencies are noted, the TFM is revised until it is acceptable.

Step 7. Evaluate Security Features User's Guide. The analysts evaluate the documentation provided to guide users in using the AIS securely, as required in the SFUG criterion shown in Figure 5.4 and discussed in section 5.5.

To ascertain whether this requirement is met, the analysts answer the following questions:

- If the system is built on an EPL product:
 - Is the evaluated SFUG included in the SFUG for the AIS?
 - Are any additional instructions or warnings for users necessary? If so, is the SFUG for the EPL product appropriately supplemented?
- If the system is not built on an EPL product or is built on an EPL product in other than its evaluated configuration:
 - Items 2 through 10 in Section 2.14 of the *Trusted Product Evaluation Questionnaire*.

If any deficiencies are noted, the documentation is revised until it is acceptable.

Step 8. Review Security Test Plan. The analysts review the plan for testing the security features of the AIS as required in the System Testing criterion shown in Figure 5.5 and discussed in section 5.6.

To ascertain whether this requirement is met, the analysts answer the following questions:

- If the system is built on an EPL product:
 - Is the vendor-provided test suite included?
 - Does the test plan address assurances that the applications do not affect the TCB?
 - Does the test plan ensure that the system is configured as specified in the EPL and TFM and in site-specific requirements and operation-concept documentation?
 - Does the test plan test all additional security-related functionality built on top of the evaluated TCB?
 - Items 4 through 8 in Section 2.11 of the *Trusted Product Evaluation Questionnaire*.
- If the system is not built on an EPL product or is built on an EPL product in other than its evaluated configuration:
 - Items 4 through 8 in Section 2.11 of the *Trusted Product Evaluation Questionnaire*.

All identified deficiencies in the test plan must be corrected.

Step 9. Review Security Test Procedures. Once the test plan is acceptable, the analysts review the test procedures to ensure that they are clear, appropriate, and complete. Although the format and media may vary (e.g., on-line, hard copy), the test procedures generally should include the following information for each test described in the test plan.

- Test Name
- Brief Description
- Requirement Being Tested
- Test Environment (equipment configuration, test programs, etc.)
- Inputs
- Expected Outputs
- Test Script

If the procedures are not acceptable, they must be corrected before proceeding.

Step 10. Conduct Security Testing. Once the test plans and procedures are acceptable, the analysts conduct the tests. In addition, they should implement at least five system-specific tests in an attempt to circumvent the security mechanisms of the system by exploiting "obvious" flaws (see section 5.6). If deficiencies are identified, the analysts identify the risks associated with the deficiencies (see section 7.2) to assess the impact to the security of the AIS and to determine whether the problems can and should be corrected. If so, the required changes are made, and the tests are rerun. If no fix is possible, the analysts determine whether the problems can be handled via procedural work-arounds. If so, the analysts document the procedures necessary to minimize the risk, and describe any residual risk that remains despite the work-around. If neither a fix nor a procedural work-around is possible, the analysts document the risk.

Step 11. Summarize Findings. The analysts summarize their findings and recommendations. The summary should include a discussion of risks identified during the assessment and measures developed to counter or lessen those risks. If the system is built on an EPL product, and the test results and additional assurances are acceptable, the analysts recommend approval. If the system appears to adequately provide the required controlled access protection mechanisms and assurances, but is not built on an EPL product or is built on an EPL product in other than its evaluated configuration, the analysts recommend interim acceptance and specify an effectiveness time period.

Chapter 7

RISK MANAGEMENT

Because absolute security is neither technically nor theoretically attainable in a multi-user system, determining whether an AIS is "secure" is essentially an exercise in identifying risks and counterbalancing those risks against protection mechanisms. So the ultimate objective of any security program is risk management.

Risk analysis is the part of risk management used to minimize risk by effectively applying security measures commensurate with the relative threats, vulnerabilities, and values of the resources to be protected. The value of a resource (e.g., AIS, data, facility) considers both its role in the organization to which it belongs and the impact that would result from its loss or unauthorized modification. Risk analysis provides a systematic way for managers to evaluate the relative costs and benefits of various security measures and to identify those that are necessary to reduce system risks to an acceptable level ("acceptable" being system specific). *Risk* is a measure of the potential for loss resulting from a *threat* and the system's *vulnerability* to that threat. A *threat* is a means by which a person (intentionally or unintentionally) or an event (natural or fabricated) can exploit a vulnerability to adversely affect the system. A *vulnerability* is a weakness that could be exploited to cause some degradation or loss of ability to perform the designated mission [10], [31], [30], [35].

To illustrate these terms, consider the case of the Trojan horse described in section 4.2. Here, the vulnerability is the DAC mechanism; the threat is a user who writes the malicious Trojan horse; and the risk is the probability that protected information will be lost as a result of the malicious program's being executed. As this example illustrates, the amount of risk is relative to the value of the information that could be compromised. If the value of the information is minimal (say, sensitive unclassified information), then the risk (probability of being exposed to unauthorized users) may be acceptable, and DAC protection will be adequate (i.e., an acceptable balance of risk versus cost is attained). However, if the value of the information is high (say, TOP SECRET), then the risk (say, the probability of the information's being exposed to

SECRET-cleared users) may not be acceptable, and more stringent security controls (i.e., mandatory access controls) may be necessary.

Risk analysis occurs throughout the system's life cycle to ensure that the security policy enforced is appropriate relative to the assumed risk. During the technical analysis described in Chapter 6, whenever a deficiency in a required security mechanism is identified, the analysts assess the risk associated with any vulnerability produced by that deficiency. When security testing is completed, the analysts assess risk on a global, system level. Three categories of risks are:

1. Risks external to the AIS.
2. Risks related to the threats that controlled access protection addresses but is not capable of countering.
3. Risks associated with shortfalls uncovered during the technical analysis.

The first two categories concern risks that the AIS organization must consider when specifying the security policy and requirements. The first category may be handled through the application of other security disciplines such as physical, communications (e.g., encryption), and operations security, or TEMPEST isolation. In specifying requirements, the AIS organization must consider the risks that these disciplines address as well as AIS risks. Category two involves specific, known limitations of controlled access protection; these limitations are discussed in section 7.1 below. The third category, addressed in section 7.2, includes risks associated with vulnerabilities that controlled access protection is intended to counter, but does not adequately do so in the AIS being analyzed.

7.1 PROTECTION LIMITATIONS

As pointed out in Chapter 2, controlled access protection (or a C2-rated product) is designed to provide security features to control sharing among mutually-trusting, cooperative users; it is not intended to provide protection sufficient to isolate separate classification levels or to counter active attempts to subvert or penetrate the system.

During risk assessment, the AIS organization identifies the levels and criticality of information to be stored and processed in the AIS, and characterizes the user environment. If the AIS will store or process information from more than one classification level or special-access compartment, and if some users will not be cleared for all levels and categories or if output needs to be accurately labeled relative to its sensitivity, then the mechanisms and assurances provided by controlled access protection are not

sufficient. Also, if the risks associated with susceptibility to Trojan horses are not acceptable, more protection may be needed.

A type of Trojan horse that has gained popular notoriety in recent years is the computer *virus*, which behaves similarly to a Trojan horse with the additional property that it attaches itself to executable programs. As with any other Trojan horse, DAC cannot prevent viruses from affecting the files to which the victims have legitimate access (even though those victims may be oblivious to the malicious actions they are triggering). However, the I&A mechanism will ensure that the only individuals given access to the system are those who can supply an authorized user identifier and can provide the evidence the system requires to authenticate their identity. Also, auditing of users' actions should serve to discourage such antisocial behavior and provides a useful tool for investigating suspicious (or malicious) behavior should it occur.

7.2 IDENTIFIED DEFICIENCIES

The analysts can determine risks associated with deficiencies identified during the technical analysis by considering the vulnerabilities that these deficiencies present. Some vulnerabilities associated with each type of deficiency the analysts may encounter during the technical analysis (see Figure 6.1) are discussed below. If the AIS is built on an EPL product, a good source of information regarding deficiencies during the product evaluation is the Evaluators' Comments section of the FER. This section may identify deficiencies that are important in the environment under consideration.

7.2.1 SYSTEM ARCHITECTURE

If the system architecture does not provide isolation and protection of the TCB, then the integrity of the TCB may be compromised, and any security mechanisms that the AIS purports to provide cannot be trusted to perform as claimed. In other words, if the system architecture is not sound, anything can happen. For example, if a user's application program can overwrite TCB code, data, or control parameters, then the AIS organization cannot rely upon the TCB's security mechanisms to work as claimed. Therefore, an AIS with "obvious" architectural vulnerabilities cannot be recommended for certification and should not be accredited. (See Reference [38], D.7, for exception conditions.)

7.2.2 IDENTIFICATION AND AUTHENTICATION

If the mechanisms for identifying and authenticating users are not sound, then the mechanisms that depend upon I&A integrity (e.g., discretionary access control and audit mechanisms) cannot be sound either. For example, if the audit mechanism uses erroneous user identities, then the data it collects are essentially useless. So even if the system provides TCB isolation, no individual accountability is possible. Therefore, an AIS with "obvious" vulnerabilities in its I&A mechanisms cannot be recommended for certification and should not be accredited. (See Reference [38], D.7, for exception conditions.)

7.2.3 DISCRETIONARY ACCESS CONTROL

The DAC mechanism may be deficient for a number of reasons (e.g., not enforced to the granularity of a single user, not enforced on all named objects controlled by the TCB), and the attendant risks will depend upon the specific deficiency and the processing environment. For example, if each user is confined to a restricted menu-driven environment, then the fact that DAC is not enforced on every object may be of little consequence. However, the lack of enforcement to the granularity of a single user may be a problem. On the other hand, if the installation's concept of operations is based on the allocation of responsibility to groups of individuals, then a group-based DAC mechanism may be adequate.

7.2.4 OBJECT REUSE

As with DAC, the object reuse mechanism may be deficient for a number of reasons. For example, the TCB may clear some storage objects (e.g., memory pages, disk blocks), but may not clear all storage objects (e.g., registers, buffers). The risk will be related to the system design, and the risk assessment should seek to determine under what conditions users may be able to see information left from a previous user's process (including the possible inclusion of this information in objects that are output from the system) and the consequences should this actually happen.

7.2.5 AUDIT

A common deficiency in the auditing mechanism is the failure of the mechanism to ensure that audit data cannot be overwritten or otherwise lost (e.g., in the case of a system crash). Usually, the effect of this deficiency can be countered with operational procedures (e.g., saving the audit trail regularly and well before it is likely to be full)

or by having the system halt when the storage space reserved for the audit trail is approaching saturation (which will result in denial of service and possible loss of some audit data). Also, AISs frequently lack good tools for audit reduction and analysis. The risks associated with these deficiencies depend upon how reliant the system is on auditing as a deterrent to malicious behavior and as a means of investigating possible misuse or abuse of the AIS.

7.2.6 SYSTEM INTEGRITY

The capability to assure that the TCB hardware has been correctly initialized and can be periodically validated is critical. Unless a mechanism exists for gaining this assurance, the requirement for TCB isolation cannot be met, and the vulnerability is the same as if the System Architecture criterion had not been met. If no system-integrity mechanism is provided, a procedure for assuring the integrity of the TCB hardware at system initialization and for validating correct operation should be implemented.

Bibliography

- [1] *A Guide to Understanding Audit in Trusted Systems*, Report No. NCSC-TG-001, National Computer Security Center, Ft. George G. Meade, MD, 1 June 1988.
- [2] *A Guide to Understanding Configuration Management in Trusted Systems*, Report No. NCSC-TG-006, National Computer Security Center, Ft. George G. Meade, MD, 28 March 1988.
- [3] *A Guide to Understanding Data Remanence in Automated Information Systems*, Report No. NCSC-TG-025, National Computer Security Center, Ft. George G. Meade, MD, June 1991.
- [4] *A Guide to Understanding Design Documentation in Trusted Systems*, Report No. NCSC-TG-007, National Computer Security Center, Ft. George G. Meade, MD, 2 October 1988.
- [5] *A Guide to Understanding Discretionary Access Control in Trusted Systems*, Report No. NCSC-TG-003, National Computer Security Center, Ft. George G. Meade, MD, 30 September 1987.
- [6] *A Guide to Understanding Identification and Authentication in Trusted Systems*, Report No. NCSC-TG-017, National Computer Security Center, Ft. George G. Meade, MD, September 1991.
- [7] *A Guide to Understanding Object Reuse in Trusted Systems*, Report No. NCSC-TG-018, National Computer Security Center, Ft. George G. Meade, MD, 1 July 1991.
- [8] *A Guide to Writing the Security Features User's Guide for Trusted Systems*, Report No. NCSC-TG-026, National Computer Security Center, Ft. George G. Meade, MD, September 1991.
- [9] Chokhani, S. "System Architecture Requirements in Trusted Computing Bases," MITRE Working Paper No. 89W00262, August 1989.
- [10] *Computer Security Policy*, AFR 205-16, Department of the Air Force, 28 April 1989.

- [11] *Computer Security Requirements: Guidance for Applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments*, Report No. CSC-STD-003-85, Department of Defense, June 25, 1985.
- [12] *Configuration Management Practices for Systems, Equipment, Munitions, and Computer Programs*, MIL-STD-483, Department of Defense, 4 June 1985.
- [13] *Department of Defense Password Management Guideline*, Report No. CSC-STD-002-85, Department of Defense, 12 April 1985.
- [14] *Department of Defense Trusted Computer System Evaluation Criteria*, DOD 5200.28-STD, Department of Defense, December 1985.
- [15] *Final Evaluation Report, SKK, Incorporated, Access Control Facility 2 (ACF2)*, Report No. CSC-EPL-84/002, National Computer Security Center, Ft. George G. Meade, MD, 1984.
- [16] *Final Evaluation Report, American Telephone and Telegraph Co. (AT&T), System V/MLS*, Report No. CSC-EPL-90/003, National Computer Security Center, Ft. George G. Meade, MD, 1990.
- [17] *Final Evaluation Report, Data General Corporation, AOS/VS Revision 7.60 Running on MV/ECLIPSE Series Processors*, Report No. CSC-EPL-89/001, National Computer Security Center, Ft. George G. Meade, MD, 1989.
- [18] *Final Evaluation Report, Digital Equipment Corporation, VAX/VMS Version 4.3*, Report No. CSC-EPL-86/004, National Computer Security Center, Ft. George G. Meade, MD, 1986.
- [19] *Final Evaluation Report, Hewlett Packard Computer Systems Division, MPE V/E*, Report No. CSC -EPL-88/010, National Computer Security Center, Ft. George G. Meade, MD, 1988.
- [20] *Final Evaluation Report, Honeywell Information Systems Multics MR11.0*, Report No. CSC-EPL-85/003, 1985.
- [21] *Final Evaluation Report, International Business Machines Corporation Multiple Virtual Storage/System Product (MVS/SP)*, Report No. CSC-EPL-90/002, National Computer Security Center, Ft. George G. Meade, MD, 1990.
- [22] *Final Evaluation Report of International Business Machines Corporation MVS/XA with RACF Version 1.8*, CSC-EPL-88/003, National Computer Security Center, Ft. George G. Meade, MD, 15 June 1988.
- [23] *Final Evaluation Report, International Business Machines Corporation, Resource Access Control Facility (RACF)*, Report No. CSC-EPL-84/001, National Computer Security Center, Ft. George G. Meade, MD, 23 July 1984.

- [24] *Final Evaluation Report, Prime Computer Corporation, Primos revision 21.0.1DODC2A*, Report No. CSC-EPL-88/009, 1988.
- [25] *Final Evaluation Report, Trusted Information Systems, Trusted XENIX*, Report No. CSC-EPL-91/003, National Computer Security Center, Ft. George G. Meade, MD, 1990.
- [26] *Final Evaluation Report, Unisys Corporation A Series MCP/AS*, Report No. CSC-EPL-18/003, National Computer Security Center, Ft. George G. Meade, MD, 1987.
- [27] *Final Evaluation Report, Unisys OS 1100 Security Release 1*, Report No. CSC-EPL-89/004, National Computer Security Center, Ft. George G. Meade, MD, 1989.
- [28] *Final Evaluation Report, Wang Laboratories, Inc., SVS/OS CAP 1.0*, Report No. CSC-EPL-90/004, National Computer Security Center, Ft. George G. Meade, MD, 1990.
- [29] *Glossary of Computer Security Terms*, NCSC-TG-004, National Computer Security Center, Ft. George G. Meade, MD, 21 October 1988.
- [30] *Guideline for Computer Security Certification and Accreditation*, FIPS PUB 140, 17 September 1983.
- [31] *Guidelines for Automatic Data Processing Physical Security and Risk Management*, FIPS PUB 31, June 1974.
- [32] *Guidelines for Writing Trusted Facility Manuals*, Report No. NCSC-TG-016, National Computer Security Center, Ft. George G. Meade, MD, date TBD.
- [33] *Information System Security Officer Guideline*, Report No. NCSC-TG-0, National Computer Security Center, Ft. George G. Meade, MD, June 1991.
- [34] *Information Systems Security Products and Services Catalogue*, National Security Agency, Ft. George G. Meade, MD, published quarterly.
- [35] *Management of Federal Information Resources*, 12 December 1985.
- [36] *National Policy on Controlled Access Protection*, NTISSP No. 200, 15 July 1987.
- [37] Neumann, Peter G. *On the Design of Dependable Computer Systems for Critical Applications*, SRI International, Computer Science Laboratory, SRI-CSL-90-10, October 1990.
- [38] *Security Requirements for Automated Information Systems (AISs)*, Department of Defense Directive 5200.28, March 21, 1988.

- [39] *Trusted Database Management System Interpretation of the Trusted Computer System Evaluation Criteria*, NCSC-TG-021, National Computer Security Center, Ft. George G. Meade, MD, April 1991.
- [40] *Trusted Product Evaluation Questionnaire*, Report No. NCSC-TG-019, National Computer Security Center, Ft. George G. Meade, MD, 2 May 1992.
- [41] *Trusted Product Evaluations: A Guide for Vendors*, Report No. NCSC-TG-002, National Computer Security Center, Ft. George G. Meade, MD, 22 June 1990.

Chapter 8

ACRONYMS

ACL	Access Control List
AIS	Automated Information System
CCB	Configuration Control Board
CI	Configuration Item
CM	Configuration Management
CRB	Configuration Review Board
DAA	Designated Approving Authority
DAC	Discretionary Access Control
DoD	Department of Defense
EPL	Evaluated Product List
FOCI	Foreign Ownership, Control, or Influence
FFRDC	Federally Funded Research and Development Corporation
I&A	Identification and Authentication
I/O	Input/Output
IPAR	Initial Product Assessment Report
ISSO	Information System Security Officer
MOU	Memorandum of Understanding
NCSC	National Computer Security Center
NSA	National Security Agency
NTIS	National Technical Information Service
RAMP	RAting Maintenance Phase
SFUG	Security Features User's Guide
TCB	Trusted Computing Base
TCSEC	Trusted Computer System Evaluation Criteria
TFM	Trusted Facility Manual
TPEP	Trusted Product Evaluation Program
TRB	Technical Review Board

Chapter 9

GLOSSARY

Except for "technical analysis," the following definitions are derived from a number of sources. [10] [2] [5] [29] [38] [14]

access A specific type of interaction that results in the flow of information between a subject and an object.

access control list A discretionary access control mechanism that implements an access control matrix by representing the columns as lists of users attached to the protected objects.

access control matrix A two-dimensional matrix representing users on the rows and objects on the columns. Each entry in the matrix represents the access type held by that user to that object. Access control matrices are usually sparsely populated and are represented in memory by row or by column, eliminating storage requirements for empty entries.

accountability The property that enables activities on an AIS to be traced to individuals who may then be held responsible for their actions.

accreditation The formal declaration by a Designated Approving Authority (DAA) that an AIS is approved to operate in a particular security mode, using a prescribed set of safeguards. Accreditation is the official management authorization for operation of an AIS and is based on the certification process as well as other management considerations. The accreditation statement affixes security responsibility with the DAA and shows that due care has been taken for security.

assurance A measure of confidence that the security features and architecture of an AIS accurately mediate and enforce the security policy. If the security features of an AIS are relied on to protect classified or sensitive unclassified information

and restrict user access, the features must be tested to ensure that the security policy is enforced and may not be circumvented during AIS operation.

automated information system (AIS) An assembly of computer hardware, software, and/or firmware configured to collect, create, communicate, compute, disseminate, process, store, and/or control data or information.

capability A protected identifier that both identifies the object and specifies the access rights to be allowed to the accessor who possesses the capability. Two fundamental properties of capabilities are that they may be passed from one accessor (subject) to another, and that the accessor who possesses capabilities may not alter or fabricate capabilities without the mediation of the operating system TCB.

certification The technical evaluation of an AIS's security features and other safeguards, made in support of the accreditation process, which establishes the extent to which a particular AIS design and implementation meet a set of specified security requirements.

configuration management The management of changes made to a system's hardware, software, firmware, documentation, tests, test fixtures, and test documentation throughout the development and operational life of the system.

controlled access protection The provision of security mechanisms and assurances that enforce a finely grained discretionary access control policy, making users individually accountable for their actions through login procedures, auditing of security-relevant events, and resource isolation.

dedicated security mode A mode of operation wherein all users have the clearance or authorization and need-to-know for all data handled by the AIS. If the AIS processes special access information, all users require formal access approval. In the dedicated mode, an AIS may handle a single classification level and/or category of information or a range of classification levels and/or categories.

Designated Approving Authority (DAA) The official who has the authority to decide on accepting the security safeguards prescribed for an AIS or the official who may be responsible for issuing an accreditation statement that records the decision to accept those safeguards. The DAA must be at an organizational level, and must have authority to evaluate the overall mission requirements of an AIS and to provide definitive directions to AIS developers or owners relative to the risk in the security posture of the AIS.

discretionary access control (DAC) A means of restricting access to objects based upon the identity of subjects and/or groups to which they belong. The controls are discretionary in the sense that a subject with a certain access permission is

capable of passing that permission (perhaps indirectly) on to any other subject (unless restrained by mandatory access control).

domain The set of objects that a subject has the ability to access.

Evaluated Products List (EPL) A documented inventory of equipment, hardware, software, and/or firmware that have been evaluated by the National Security Agency against the evaluation criteria found in DoD 5200.28-STD. The EPL is maintained by the NSA's National Computer Security Center (NCSC).

identification and authentication (I&A) The combination of a process that enables recognition of an entity by a system, generally by the use of unique machine-readable user names (identification) and the verification of the identity of a user, device, or other entity in a computer system, often as a prerequisite to allowing access to resources in a system (authentication).

Information System Security Officer (ISSO) The person responsible to the DAA for ensuring that security is provided for and implemented throughout the life cycle of an AIS from the beginning of the concept development phase through its design, development, operation, maintenance, and secure disposal.

mandatory access control (MAC) A means of restricting access to objects based on the sensitivity (as represented by a label) of the information contained in the objects and the formal authorization (i.e., clearance) of subjects to access information of such sensitivity.

multilevel security mode A mode of operation that allows two or more classification levels of information to be processed simultaneously within the same system when not all users have a clearance or formal access approval for all data handled by the AIS.

object A passive entity that contains or receives information. Access to an object potentially implies access to the information it contains. Examples of objects are: records, blocks, pages, segments, files, directories, directory trees, and programs, as well as bits, bytes, words, fields, processors, video displays, keyboards, clocks, printers, and network nodes.

object reuse The reassignment to some subject of a medium (e.g., page frame, disk sector, magnetic tape) that contained one or more objects. To be securely reassigned, such media must contain no residual data from the objects they previously contained.

password (1) A protected/private character string used to authenticate an identity.
(2) A discretionary access control mechanism that represents the access control matrix by row by attaching passwords to protected objects.

- profile** A discretionary access control mechanism that associates a list of protected objects with each user.
- protection bits** An incomplete attempt to represent the access control matrix by column. Implementation of protection bits include systems such as Unix, which use protection bits associated with objects instead of a list of users who may access an object.
- privileged instructions** A set of instructions (e.g., interrupt handling or special computer instructions) that control features (such as storage protection features) and that are generally executable only when the automated system is operating in the executive state.
- reference monitor concept** An access control concept that refers to an abstract machine that mediates all accesses to objects by subjects.
- risk** A combination of the likelihood that a threat shall occur, the likelihood that a threat occurrence shall result in an adverse impact, and the severity of the resulting adverse impact.
- risk analysis** An analysis of system assets and vulnerabilities to establish an expected loss from certain events based on estimated probabilities of occurrence.
- risk management** The total process of identifying, measuring, and minimizing uncertain events affecting AIS resources. This process includes risk analysis, cost benefit analysis, safeguard selection, security test and evaluation, safeguard implementation, and systems review.
- security policy** The set of laws, rules, and practices that regulate how an organization manages, protects, and distributes sensitive information.
- storage object** An object that supports both read and write accesses.
- subject** An active entity, generally in the form of a person, process, or device that causes information to flow among objects or changes the system state. Technically, a process/domain pair.
- system high security mode** A mode of operation wherein all users having access to the AIS possess a security clearance or authorization, but not necessarily a need-to-know, for all data handled by the AIS. If the AIS processes special access information, all users must have formal access approval.
- system integrity** The quality that a system has when it performs its intended function in an unimpaired manner, free from deliberate or inadvertent unauthorized manipulation of the system.

technical analysis The process of examining an AIS to ensure that it meets its functional, assurance, and documentation requirements; within the context of this guidance, specifically relative to controlled access protection (see "certification").

threat The means through which the ability or intent of a threat agent to adversely affect an automated system, facility, or operation can be manifest. Threats may be categorized and classified as: intentional or unintentional human threats; or natural or fabricated environmental threats.

Trojan horse A computer program with an apparently or actually useful function that contains additional (hidden) functions that surreptitiously exploit the legitimate authorizations of the invoking process to the detriment of security; for example, making a "blind copy" of a sensitive file for the creator of the Trojan-horse program.

Trusted Computing Base (TCB) The totality of protection mechanisms within a computer system — including hardware, firmware, and software — the combination of which is responsible for enforcing a security policy. A TCB consists of one or more components that together enforce a unified security policy over a product or system. The ability of a TCB to correctly enforce a security policy depends solely on the mechanisms within the TCB and on the correct input by system administrative personnel of parameters (e.g., a user's clearance) related to the security policy.

trusted product A product that has been evaluated and approved for inclusion on the Evaluated Products List.

user Any person who interacts directly with a computer system.

virus A self-propagating Trojan horse, composed of a mission component, a trigger component, and a self-propagating component.

vulnerability The characteristic of a system that causes it to suffer a definite degradation (inability to perform the designated mission) as a result of having been subjected to a certain level of effects in the unnatural (human-made) hostile environment. For computers, it is a weakness in automated system security procedures, administrative controls, internal controls, etc., that could be exploited to gain unauthorized access to information or disrupt critical processing.